

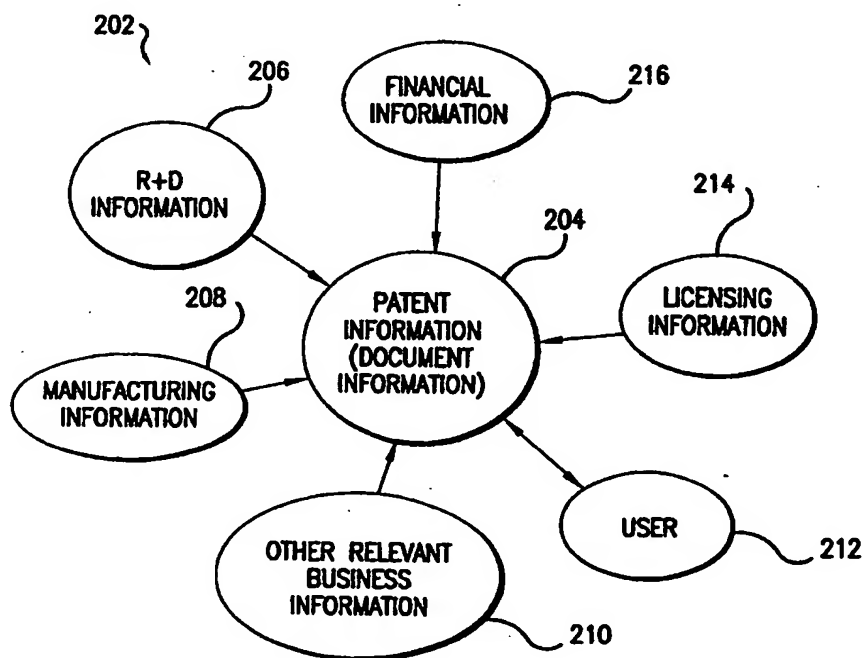
PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30		A1	(11) International Publication Number: WO 98/55945
			(43) International Publication Date: 10 December 1998 (10.12.98)
(21) International Application Number: PCT/US98/10923		(74) Agents: LEE, Michael, Q. et al.; Sterne, Kessler, Goldstein & Fox P.L.L.C., Suite 600, 1100 New York Avenue, N.W., Washington, DC 20005-3934 (US).	
(22) International Filing Date: 2 June 1998 (02.06.98)			
(30) Priority Data:		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
08/867,392 2 June 1997 (02.06.97) US			
08/921,369 29 August 1997 (29.08.97) US			
(71) Applicant: SMARTPATENTS, INC. [US/US]; 1975 Landings Drive, Mountain View, CA 94043 (US).			
(72) Inventors: RIVETTE, Kevin, G.; 2165 Waverley Street, Palo Alto, CA 94303 (US). RAPPAPORT, Irving, S.; 1500 Edgewood Drive, Palo Alto, CA 94303 (US). HOHMANN, Luke; 306 Windmill Park Lane, Mountain View, CA 94043 (US). PUGLIA, David; 17429 East Vineland Avenue, Los Gatos, CA 95030 (US). GORETSKY, David; 272 Waverly Street, Sunnyvale, CA 94086 (US). JACKSON, Adam; 1063 Morse Avenue #7-107, Sunnyvale, CA 94089 (US). RABB, Charles, Jr.; 730 E. Evelyn #638, Sunnyvale, CA 94086 (US). SMITH, David, W.; 3 Morning Sun Court, Mountain View, CA 94043 (US). PARK, Brian; 4029 Park Boulevard, Palo Alto, CA 94306 (US). THORNTHWAITE, Warren; 147 Hedge Road, Menlo Park, CA 94025 (US). NAVARRETE, Jorge, A.; 160 Hedge Road, Menlo Park, CA 94025 (US).			

Published*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.***(54) Title:** SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR PATENT-CENTRIC AND GROUP-ORIENTED DATA PROCESSING, INCLUDING USING HYPERBOLIC TREES TO VISUALIZE DATA**(57) Abstract**

A system, method, and computer program product for processing data are described herein. The system maintains first databases of patents, and second databases of non-patent information of interest to a corporate entity. The system also maintains one or more groups. Each of the groups comprises any number of the patents from the first databases. The system, upon receiving appropriate operator commands, automatically processes the patents in one of the groups in conjunction with non-patent information from the second databases. Accordingly, the system performs patent-centric and group-oriented processing of data. A group can also include any number of non-patent documents. The groups may be product based, person based, corporate entity based, or user-defined. Other types of groups are also covered, such as temporary groups. The processing automatically performed by the system relates to (but is not limited to) patent mapping, document mapping, patent citation (both forward and backward), patent aging, patent bracketing/clustering (both forward and backward), inventor patent count, inventor employment information, patent claim tree analysis, and finance. Other functions and capabilities are also covered, including the ability to utilize hyperbolic trees to visualize data generated by the system, method, and computer program product.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR
PATENT-CENTRIC AND GROUP-ORIENTED
DATA PROCESSING, INCLUDING USING HYPERBOLIC TREES TO VISUALIZE DATA

Background of the Invention

5 *Field of the Invention*

The present invention is generally related to tools for data processing, and more particularly related to tools for patent-centric and group-oriented data processing. These tools comprise diverse capabilities for data presentation and processing, including data presentation and processing using hyperbolic trees.

10 *Related Art*

Patents are becoming more and more important to a business's success, especially in today's global economy. Patents can be viewed as a new type of currency in this global economy because they grant the holder with a right to exclude others from making, using, or selling the patented technology. In some
15 industries, product turnover is fairly rapid. However, core technology, product features, and markets change at a much slower rate. Accordingly, even in fast-moving industries, patents which cover core technology are very valuable at protecting a company's research and development investment for an extended period of time.

20 Patents are also valuable as revenue generators. In 1993, for example, the revenue generated from patents by U.S. companies was over \$60 billion. Fred Warshofsky, *The Patent Wars*, John Wiley & Sons, Inc., New York, 1994. These patent revenue dollars are rising each year.

25 Patents are further valuable because they collectively represent a vast technological database. Much of this database is only available as issued patents (i.e., it is not released in any other form). According to Larry Kahaner's book,

-2-

Competitive Intelligence, Simon & Schuster, 1996, "More than 75 percent of the information contained in U.S. patents is never released anywhere else."

If corporations searched this database before developing and releasing new products they might be able to avoid costly patent infringement litigation. Often, however, corporations do not conduct such patent searches. One significant reason for this is the difficulty in identifying relevant patents, and the difficulty in analyzing patents. Computerized search tools are becoming available to the public, such as web sites on the Internet, that can be used to conduct patent searches. Many companies and practitioners are reluctant to use such tools, however, due to the concern that their highly sensitive patent searches will not be maintained in confidence when using such tools.

More and more corporations are recognizing the value of patents. The number of patents applied for and issued to U.S. companies is increasing every year, especially in fast moving industries such as computer software and biotechnology. Many international companies have also recognized the value of patents. In fact, foreign companies regularly rank among the leaders in issued U.S. patents.

Of course, not all patents are as valuable to the patent owner or patent licensees as others. Some owned or licensed patents provide little or no value to the corporate entity. These patents become a drain on corporate resources, both in obtaining the patents, paying maintenance fees, and paying license fees. It is difficult for corporations to assess the value of their patents because automated tools for patent analysis do not exist.

Yet, for all the heightened awareness being paid to patents in some quarters, patents remain one of the most underutilized assets in a company's portfolio. This is due, at least in significant part, to the fact that patent analysis, whether for purposes of licensing, infringement, enforcement, freedom to operate, technical research, product development, etc., is a very difficult, tedious, time

consuming, and expensive task, particularly when performed with paper copies of patents.

Software providers have been slow in developing software tools for aiding in the patent analysis process. As a result, there are few automated tools for patent analysis currently available. There are software tools available for managing corporate patent prosecution and payment of maintenance fees, such as products from Master Data Corporation. The patent analysis capabilities of these tools are limited. These tools, for example, cannot be used to facilitate the analysis and development of business strategies to increase corporate shareholder value through the strategic and tactical use of patents.

A number of patent searching tools are available, such as the United States Patent and Trademark Office (USPTO) Automated Patent System (APS), and the on-line search services offered by Lexis and Westlaw. Other providers of patent information and patent search tools include Derwent, MicroPatent, Questel, Corporate Intelligence, STN, IFI/Plenum, The Shadow Patent Office (EDS), IBM, and CAS. These tools are not analysis tools. Instead, they are search tools. These tools enable a user to identify patents that satisfy a specified key word search criteria. In essence, these tools provide the user with the ability to possibly find "the needle-in-the-haystack." However, these tools have limited, if any, automated functions to aid a user in analyzing the patents, whether the company's own patents or those of competitors, for the purpose of making tactical and strategic business decisions based on the patents.

SmartPatents Inc. (SPI) of Mountain View, CA, provides electronic tools for analyzing patents. These tools, collectively called the SmartPatent Workbench, are very useful for analyzing patents. With the SmartPatent Workbench, a user can view the text and image of a patent, conduct text searches in the patent, copy and paste portions of the patent to other documents, build a case of patents, annotate the case and the patents in the case, import and export patents and cases, etc. The SmartPatent Workbench is commercially available

from SPI, and is described in a number of publicly available documents, such as U.S. Patent No. 5,623,679 and U.S. Patent No. 5,623,681, incorporated by reference herein.

5 The SmartPatent Workbench is a patent analysis tool. The SmartPatent Workbench is primarily designed to assist a user in working with a single patent or a small collection of patents at a time. However, there are many instances when it would be very beneficial to be able to automatically and simultaneously analyze, correlate, or otherwise process multiple patents.

10 For example, in some instances it would be beneficial to automatically analyze the inventorship of a collection of patents. More particularly, it would be beneficial to identify the persons who are named most frequently on a collection of patents. It would be very useful if this task could be performed automatically. However, no existing software tools can perform this task automatically.

15 For the most part, existing patent-related tools can process only the information contained in patents. (It is noted, however, that the SmartPatent Workbench has functions to annotate patents with any information, whether or not patent related, and has additional functions to search within annotations.) These tools do not have functions for correlating, analyzing, and otherwise processing patent-related information with non-patent related information, including but not
20 limited to corporate operational data, financial information, production information, human resources information, and other types of corporate information. Such non-patent information is critically important when evaluating the full strategic and tactical value and applicability of any given patent, or developing a corporate patent business strategy for gaining competitive advantage
25 and increasing shareholder value based on patents.

Consider, for example, FIG. 1. A typical corporation 102 includes a research and development (R&D) department 104, a finance department 112, a manufacturing department 108, and a legal department 116 (that includes a licensing department 122 and a patent department 124). In the course of

performing their respective duties, these departments generate, collect, and maintain information, such as R&D information 106, financial information 114, manufacturing information 110 (such as bill of material information), licensing information 118, and patent information 120 (that includes the patents obtained by the company, and perhaps patents obtained by competitors).

A business analyst 126 may be assigned the job of evaluating the value of the corporation's patent portfolio (represented as part of the patent information 120). In order to fully and accurately analyze the value and applicability of the corporation's patent portfolio, the analyst 126 should ideally take into account non-patent information, such as R&D information 106, financial information 114, manufacturing information 110, and licensing information 118.

For example, a patent's value may be linked to whether it covers technology that the corporation is currently using, or that the corporation may use in the future. Thus, an analysis of the patent should include an analysis of and correlation with manufacturing information 110 and R&D information 106. Also, a patent's value may be linked to whether it has generated licensing revenue. Thus, an analysis of the patent should include an analysis of and correlation with licensing information 118. Further, a patent's value may be linked to the degree of success of the corporation's commercial products that correspond to the patent (i.e., the commercial embodiments of the patented technology). Thus, an analysis of the patent should include an analysis of and correlation with financial information 114.

The processing described above, however, is usually not done (or it is done in an ad hoc, unorganized, incomplete, inefficient, and/or ineffective manner) because it is difficult or, in many cases, impossible to manually collect, organize, correlate, and process all of the information pertinent to the patents under study. Often times, it is a difficult or even impossible task to simply identify the relevant patents. Accordingly, it would be very beneficial to have automated tools that automatically process patent-related information and non-patent related

information for making corporate business decisions. Existing patent-related tools do not have this capability.

Summary of the Invention

5 Briefly stated, the present invention is directed to a system, method, and computer program product for processing data. The present invention maintains first databases of patents, and second databases of non-patent information of interest to a corporate entity.

10 The present invention also maintains one or more groups. Each of the groups comprises any number of patents from the first databases. The present invention, upon receiving appropriate operator commands, automatically processes the patents in one or more of the groups in conjunction with non-patent information from the second databases. Accordingly, the present invention performs patent-centric and group-oriented processing of data.

A group can also include any number of non-patent documents.

15 The groups may be defined by the business practices of the corporation and could include groupings that are product based, person based, corporate entity based, or user-defined. Other types of groups also fall within the scope of the invention. For example, the invention supports temporary groups that are automatically generated in the course of the automatic processing performed by the invention.

20 The processing automatically performed by the invention relates to (but is not limited to) patent mapping, document mapping, document/patent citation (both forward and backward), document/patent aging, patent bracketing/clustering (both forward and backward), inventor patent count, inventor employment information, and finance. Other functions also fall within the scope of the invention.

The present invention includes the ability to display data in a wide range of formats, including the ability to display and process data using hyperbolic trees.

Further features and advantages of the invention, as well as the structure and operation of various embodiments of the invention, are described in detail below with reference to the accompanying drawings. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements. The drawing in which an element first appears is indicated by the leftmost digit(s) in the corresponding reference number.

Brief Description of the Figures

The present invention will be described with reference to the accompanying drawings, wherein:

FIG. 1 represents the generation and maintenance of documents in a conventional corporate entity;

FIG. 2 illustrates the document-centric and patent-centric operation of the present invention;

FIG. 3 is a block diagram of a system according to a preferred embodiment of the present invention;

FIG. 4 is a block diagram of an enterprise server according to a preferred embodiment of the present invention;

FIG. 5 illustrates a potential deployment of the enterprise server of FIG. 4;

FIG. 6 is a block diagram of the databases of the present invention;

FIG. 7 is a block diagram of a network client (and potentially a web client) according to an embodiment of the invention;

FIG. 8 is a block diagram of a web server according to an embodiment of the invention;

FIG. 9 is a block diagram and a data transfer diagram illustrating the searching features of the present invention;

FIG. 10 is a block diagram of the analysis modules which form a part of the enterprise server of FIG. 4;

5 FIG. 11 is a block diagram of a computer useful for implementing components of the invention;

FIG. 12A illustrates the orientation of FIGS. 12B-12M relative to one another;

10 FIGS. 12B-12M illustrates the tables and attributes in the databases of FIG. 6 according to an embodiment of the invention;

FIGS. 13-17 illustrate example document databases;

FIG. 18 illustrates an example display format depicting the hierarchical organization of groups according to the present invention;

FIGS. 19-21 illustrates example group tables;

15 FIGS. 22 and 23A illustrate example bill of materials (BOM) data structures (also called BOM structures, or BOMs);

FIG. 23B, when considered in conjunction with FIG. 23A, illustrate the concept of shared groups;

FIGS. 24-26 illustrate example BOM groups;

20 FIGS. 27-31 illustrate example security tables;

FIG. 32 illustrates an example corporate organizational structure;

FIGS. 33-36 illustrate example corporate entity databases;

FIG. 37 illustrates an example person table;

FIG. 38 illustrates an example employee table;

25 FIG. 39 illustrates an example validated inventor table;

FIGS. 40-43, 44A and 44B illustrate example patents used to describe the patent bibliographic databases;

FIG. 45 is a dataflow diagram illustrating a generic extract and load operation;

FIG. 46 is a dataflow diagram illustrating an exemplary extract and load process for the patent bibliographic databases;

FIG. 47 is a dataflow diagram illustrating an exemplary extract and load process for the BOM databases;

5 FIG. 48 illustrates an alternative process for obtaining corporate BOM data;

FIG. 49 is a dataflow diagram representing an exemplary process for extract and load of the person databases and the employee databases;

10 FIG. 50 is a dataflow diagram illustrating an exemplary process for extract and load of the validated inventor table;

FIG. 51 is a dataflow diagram illustrating an exemplary process for extract and load of the corporate entity databases;

FIG. 52 is a dataflow diagram illustrating an exemplary process for extract and load of other corporate entity databases;

15 FIGS. 53-57 illustrate example user interface display formats pertinent to the searching features of the present invention;

FIG. 58 is an example user interface display format pertinent to display of group information;

FIGS. 59-60 are examples of patent mapping display formats;

20 FIGS. 61-65 are examples of patent citation report display formats;

FIGS. 66-70 are examples of patent aging display formats;

FIGS. 71-73 are examples of patent clustering/bracketing display formats;

FIGS. 74-77 are examples of inventor patent count display formats;

FIGS. 78-80 are examples of employment information display formats;

25 FIG. 81 illustrates the interaction between the enterprise server and a client;

FIG. 82 illustrates the interaction between the enterprise server and a network client;

FIG. 83 illustrates the interaction between the enterprise server and a web client;

FIG. 84 is a flowchart depicting the operation of the patent mapping module according to the embodiment of the invention;

5 FIG. 85 is a flowchart depicting the operation of the patent/document mapping module according to an embodiment of the invention;

FIG. 86 is a flowchart depicting the operation of the patent citation module when conducting a backward patent citation search according to an embodiment of the invention;

10 FIG. 87 is a flowchart depicting the operation of the patent citation module when performing a forward patent citation search according to an embodiment of the invention;

FIGS. 88A and 88B collectively illustrate a flowchart representing the operation of the patent aging module according to an embodiment of the invention;

15 FIG. 89 is a flowchart representing the operation of the patent bracketing/clustering module when performing a backward patent bracketing/clustering function according to an embodiment of the invention;

FIG. 90 is a flowchart illustrating the operation of the patent bracketing/clustering module when performing a forward patent bracketing/clustering function according to an embodiment of the invention;

20 FIG. 91 is a flowchart depicting the operation of the inventor patent count module according to an embodiment of the invention;

FIG. 92 is a flowchart depicting the operation of the inventor employment information module according to an embodiment of the invention;

25 FIG. 93 is a flowchart depicting the operation of the importing patent data module according to an embodiment of the invention;

FIG. 94 is a flowchart depicting the operation of the exporting patent data module according to an embodiment of the invention;

FIG. 94 is a flowchart depicting the operation of the exporting patent data module according to an embodiment of the invention;

FIG. 95 is a flowchart representative of a generic extract and load process according to an embodiment of the invention;

FIG. 96 is a flowchart of a extract and load process for the patent bibliographic databases;

5 FIG. 97 is a flowchart of a extract and load process for the BOM databases;

FIG. 98 is a flowchart of a extract and load process for an employee databases;

10 FIG. 99 is a flowchart of a extract and load process for the validated inventor databases;

FIG. 100 is an extract and load flowchart for the corporate entity databases;

FIG. 101 is a flowchart representative of the interaction between a client and the enterprise server;

15 FIG. 102 is a flowchart representative of a patent mapping and mining process;

FIG. 103 is a flowchart representative of a situation assessment process;

FIG. 104 is a flowchart representative of a competitive analysis process;

20 FIG. 105 is a flowchart representative of a clustering and/or bracketing process;

FIG. 106 is a flowchart representative of an inventor analysis process;

FIG. 107 is a flowchart representative of a financial analysis process;

FIG. 108 is a flowchart representative of a strategic planning process;

25 FIG. 109 is a flowchart representative of an example methodology process involving patent mapping and mining, situation assessment, and strategic planning process;

FIG. 110 is a flowchart depicting the operation of the security module;

FIG. 111 is an example display format showing the display of patent text in a first window and notes in a second window;

-12-

FIG. 112 is an example display format showing the display of patent text in a first window and patent image in a second window;

FIG. 113 illustrates a block diagram of the virtual patent system of the present invention;

5 FIG. 114 is a architecture block diagram of the network client (and in some embodiments the web client);

FIG. 115 is used to describe a generic group import function of the present invention;

FIG. 116 is an example user login screen shot;

10 FIGS. 117 and 118 represent an example console screen shot;

FIGS. 119 and 120 are screen shots for creating a new group;

FIGS. 121 and 122 are example screen shots for searching through the databases;

15 FIGS. 123 and 124 are example screen shots for displaying text and images of documents;

FIG. 125 is an example screen shot for creating a document note;

FIGS. 126 and 127 are example screen shots for editing group properties;

FIGS. 128 and 129 are example screen shots for invoking patent-centric and group-oriented functions;

20 FIG. 130 is an example screen shot for adding a document to a group;

FIG. 131 is an example screen shot for importing data;

FIG. 132 is an example screen shot for exporting data;

FIG. 133 is another example console screen shot;

FIG. 134 is an example screen shot for creating a group note;

25 FIGS. 135-137 illustrate example tools bars from the console screen display;

FIG. 138 illustrates a search hierarchy used to describe the searching algorithm according to a preferred embodiment of the present invention;

FIG. 139 is a flowchart depicting the operation of the present invention when performing searches according to an embodiment of the invention;

FIG. 140 illustrates an example Patent Search screen according to an embodiment of the invention;

5 FIGS. 141-143 illustrate example Search Result screens according to an embodiment of the invention;

FIG. 144 illustrates an example display screen that shows bibliographic and abstract information on a document that is not stored in the repository;

10 FIGS. 145A, 145B, and 145C illustrate an example display screen that shows information on a document that is stored in the repository;

FIG. 146 illustrates an example display screen used to illustrate the hyperlinking capabilities of the present invention;

FIG. 147 illustrates an example "Patents In Repository" screen;

15 FIG. 148 illustrates an example display screen corresponding to the Skim Images function of the present invention;

FIG. 149 is a flowchart depicting a demand paging algorithm according to an embodiment of the invention;

FIG. 150 illustrates a URL message format;

20 FIG. 151 illustrates the commands that are transferred between a browser in the web client and the Enterprise server;

FIG. 152 illustrates the interaction between the browser in a web client and the Enterprise server;

FIG. 153 illustrates a stacked folder icon used to represent shared groups;

FIG. 154 illustrates an example console used to describe shared groups;

25 FIG. 155 illustrates an example console used to describe temporary groups;

FIG. 156 illustrates a group links tab that lists a group's links in the group hierarchy;

FIGS. 157-160 are flowcharts representing the operation of the patent citation tree function when performed by a network client interacting with the enterprise server;

5 FIG. 161 is an example console used to illustrate the operation of the patent citation tree function;

FIG. 162 is an example drop-down menu used to illustrate the manner in which an operator selects the citation analysis function;

FIG. 163 is an example dialog box used to indicate how an operator defines a citation analysis command;

10 FIG. 164 illustrates an example patent citation tree;

FIG. 165 illustrates an example display that is generated when an operator selects a patent represented in the patent citation tree of FIG. 164;

15 FIGS. 166 and 167 are flowcharts representing the operation of the patent citation tree function when performed by a web client interacting with the enterprise server via the web server;

FIGS. 168-170 are flowcharts illustrating the operation of the patent claims tree function;

FIG. 171 illustrates an example patent claims tree;

20 FIGS. 172 and 173 illustrate example displays which are presented when the operator selects a claim represented in the patent claims tree of FIG. 171;

FIGS. 174 and 175 are additional patent citation visualizations according to embodiments of the invention;

FIG. 176 is a flowchart representing additional operation related to the patent citation tree function;

25 FIGS. 177 and 178 illustrate example hyperbolic trees;

FIG. 179 represents the mapping from a graph to a tree;

FIG. 180 represents an example parent/child table;

FIG. 181 illustrates a citation analysis graph corresponding to the patent/child table of FIG. 180;

-15-

FIG. 182 illustrates an example patent bibliographic information table;
 FIG. 183 illustrates an example tree corresponding to the citation analysis
 graph of FIG. 181;
 FIG. 184 illustrates an example claims dependency graph;
 5 FIG. 185 illustrates an example claims dependency tree corresponding to
 the claims dependency graph of FIG. 184; and
 FIG. 186 illustrates a web client in greater detail.

In the following text, reference is sometimes made to existing U.S. patents.
 Also, some of the figures reference or illustrate existing U.S. patents. For
 10 illustrative purposes, information from and/or about these patents has sometimes
 been modified or created in order to support the particular examples being
 discussed. Accordingly, the information provided herein about these existing U.S.
 patents should be considered to be fictional unless verified through comparison
 with copies of the actual U.S. patents that are available from the U.S. Patent and
 15 Trademark Office.

Detailed Description of the Preferred Embodiments

Table of Contents

Overview of the Invention

Components of the Invention

20	Customer Corporate Entity
	Databases
	Document Databases
	Document Bibliographic Databases
	Patent Bibliographic Databases
25	Other Document Bibliographic Databases
	Notes Database
	Groups Databases

-16-

Predefined Groups Databases

Bill of Materials (BOM) Databases

Corporate Entity Databases

Inventor Databases (and Employees and Person
Databases)

User-Defined Group Databases

Financial Databases

Security Database

Enterprise Server

Document Storage and Retrieval Module

Notes Module

Searching Module

Automatic Searches Related to Groups

Searching Algorithm

Grouping Module

Analysis Modules

Server Administration Module

Server Configuration Module

Command Dispatch Module

Clients

Network Clients

Web Clients

Enterprise Server API (Application Programming Interface)

Commands Processed by the Server Administration

Module 418

Commands Processed by the Document Storage and
Retrieval Module 408

Commands Processed by the Grouping Module 412

Commands Processed by the Notes Module 414

-17-

Commands Processed by the Analysis Modules 416

Client/Server Interaction

Patent-Centric URL Commands

Translation

5 Client Architecture

Databases

Document Bibliographic Databases

Group Databases

User Defined Groups

10 Predefined Group Databases

Bill of Materials (BOM) Databases

Corporate Entity Databases

Inventor, Employee, and Person Databases

Financial Databases

15 Security Databases

Enterprise Server and Client Functional Modules

Patent Mapping Module

Patent Citation Module

Patent Aging Module

20 Patent Clustering and Bracketing Module

Financial Module

Inventor Patent Count Module

Inventor Employment Information Module

Exporting Patent Data Module

25 Importing Patent Data Module

Methodology Embodiments

Patent Mapping and Mining

Situation Assessment

Competitive Analysis

-18-

	Clustering and/or Bracketing
	Inventor Analysis
	Financial Analysis
	Strategic Planning
5	Integrated Methodology Embodiment
	User Interface
	User Login
	Console
	Console Tool Bars
10	Creating a New Group
	Editing Group Properties
	Shared Groups
	Invoking Patent-Centric and Group-Oriented Analysis Functions
	Adding Documents to a Group
15	Adding a Document Note
	Adding a Group Note
	Searching
	Web Searching
	Importing Data
20	Exporting Data
	Data Presenting and Processing Using Hyperbolic Trees
	General Description of Hyperbolic Trees
	Patent Citation Tree
	Patent Citation Tree (Network Client)
25	Patent Citation Tree (Web Client)
	Additional Patent Citation Visualizations
	Patent Claims Tree
	Conclusion

Overview of the Invention

The present invention is directed to a system, components of the system, a method, components of the method, and a computer program product for patent-centric and group-oriented data processing. Such processing includes, but is not limited to, reporting, analyzing, and planning.

The present invention is intended to aid a corporate entity in developing business-related strategies, plans, and actions. Accordingly, the present invention is also referred to herein as a business decision system and method.

FIG. 2 is a conceptual representation of the invention. The present invention processes patent information 204, which is herein defined to include (but not limited to) U.S. and non-U.S. patents (text and/or images) and post issuance documents (such as Certificates of Correction), and patent-related information, which includes information about patents (herein called patent bibliographic information). Accordingly, the processing performed by the invention is said to be "patent-centric" or "patent-specific."

More generally, the present invention processes any documents, some of which are related to patents, and others which are unrelated to patents. These documents are preferably of interest to a business entity, and include contracts, licenses, leases, notes, commercial papers, other legal and/or financial papers, etc., as well as patents.

For illustrative purposes, the invention is often described herein with respect to patents. However, it should be understood that the invention is also applicable to all types of documents, and the structures, functions, and operations described herein are applicable to all types of documents, whether patent or non-patent.

The present invention also processes other information, preferably business-related information, including (but not limited to) research and development (R&D) information 206, financial information 216, patent licensing

-20-

information 214, manufacturing information 208, and other relevant business information 210 (which may, for example, include human resources information). This other information is generally called non-patent information (since it includes documents other than patents and may further include information from operational and non-operational corporate databases).

The present invention is adapted to maintain and process massive amounts of documents (several hundred thousand or more). It is often necessary to maintain and process this large number of documents in order to develop strategic, patent-related business plans for the customer.

According to the present invention, processing of the patent information 204 can be conducted either with or without consideration of any of the other information 206, 216, 214, 210, 208.

For example, a user 212 (who may be a business analyst) may be assigned the job of evaluating the value of the corporation's patent portfolio (represented as part of the patent information 204). In order to fully analyze the value and applicability of the corporation's patent portfolio, the user 212 must take into account other information, such as R&D information 206, financial information 216, manufacturing information 208, and licensing information 214, for both the corporation and its competitors.

For example, a patent's value may be linked to whether it covers technology that the corporation is currently using, or that the corporation may use in the future. For this and other purposes, the present invention includes functions for automatically analyzing the patent information 204 in conjunction with manufacturing information 208 and/or R&D information 206. Also, a patent's value may be linked to whether it has generated licensing revenue. For this and other purposes, the present invention includes functions for automatically analyzing the patent information 204 in conjunction with the licensing information 214. Further, a patent's value may be linked to the degree of success of the corporation's commercial products related to the patent

-21-

(i.e., the commercial embodiments of the patented technology). For this and other purposes, the present invention includes functions for automatically analyzing the patent information 204 in conjunction with the financial information 216.

5 The invention could also be used to determine the value of a corporate entity's patent portfolio for purposes of a merger or acquisition. The invention could also be used in a merger or acquisition context to determine a corporate entity's business direction. For example, if Company A is interested in acquiring Company B, Company A could use the invention to categorize all of Company B's
10 patents into groups. The nature of these groups would be an indication of the types of work that Company B is involved in. Other uses of the invention are described below. Further uses of the invention will be apparent to persons skilled in the relevant art(s) based on the discussion contained herein.

15 The present invention is group enabled. According to the present invention, a group is a data structure that includes a collection of patents. The patents in a group typically follow a common theme or characteristic (although this is not a mandatory requirement of groups). For example, a first group may include patents that map to a product being manufactured and sold by a company. A second group may include patents that map to a product or product feature
20 being considered for future manufacture and sale by a company. A third group may include patents owned by a corporate entity. A fourth group may include patents each having a particular person named as an inventor. A fifth group may include patents owned by a competitor. A sixth group may include patents related to a research project. A seventh group may include licensed patents. An eighth
25 group may include patents and/or non-patent documents related to a litigation in which the customer is involved or has an interest (such a group is also herein called a case). A ninth group may include patents and other documents arbitrarily selected by a customer.

-22-

The present invention is capable of automatically processing the patents in a group, or the patents in multiple groups (alternatively, the invention can automatically process a single patent). Accordingly, the present invention is said to support "group-oriented" data processing.

5 Being able to automatically process information on a group basis is a very important feature of the invention, and proves to be very valuable and useful. Consider the above example of FIG. 2, where the user 212 has the task of evaluating the value of the corporation's patent portfolio. Suppose that the corporation has two products on the market, Product A and Product B. Product
10 A generated \$10 million in revenue, and Product B generated \$30 million in revenue. The corporation has 5 patents that map to Product A, and 3 patents that map to Product B. If the user 212 analyzes this data without regard to groups, then the user 212 will find that the corporation's revenue per patent is \$5 million. That is, for every \$5 million in revenue, the corporation obtains
15 a patent. Suppose that a relevant industry benchmark indicates that a company should obtain a patent for every \$6 million of revenue. According to this scenario, the user 212 will conclude that the corporation is potentially seeking greater patent protection than the industry benchmark with respect to its technology.

20 Consider, now, the scenario where the user 212 analyzes the data with regard to groups, in this case a first group composed of patents that map to Product A, and a second group composed of patents that map to Product B. The user 212 will find that corporation's revenue per patent is \$2 million for the first group (i.e., patents that map to Product A), and \$10 million for the
25 second group (i.e., patents that map to Product B). According to this scenario, the user 212 will conclude that the corporation is potentially devoting too much of its patent-related resources with respect to its technology related to Product A (it is "overpatenting" technology related to Product A), and potentially

devoting too little of its patent-related resources with respect to its technology related to Product B (it is "underpatenting" technology related to Product B).

5 In addition, an analysis of the patents relative to a product may indicate that the core features or technology of the product are not patented and, thus, could be freely and legally copied by a competitor. This could adversely affect the product's price floor and revenue stream. With this information in hand, the company could then take steps to more comprehensively patent its technology (or make a conscious and knowledgeable decision to not seek further patent protection). Without group-oriented processing of the patents related to
10 the product, this information is unavailable. Without this information, the company is more likely to make unwise and costly business decisions.

As indicated by the above example, group-oriented processing yields information on a scale whose granularity is defined by the definition of the group. The information produced by group-oriented processing is specific to
15 the patents in the group. Accordingly, as with the above example, group-oriented processing is often more useful and more illuminating than non-group-processing.

Also, the invention supports hierarchically structured groups. The invention, in performing a function requested by the operator, may identify a
20 particular group. Such identification of this group may yield very useful information, as apparent from the above example. This group, however, may have a number of parent and/or child groups. The operator may be able to uncover additional useful data by viewing, analyzing, and/or processing these parent and child groups, either with or without the original group.

25 Accordingly, the invention supports and facilitates "data drilling" and/or "data mining."

As noted above, according to the present invention, processing of the patent information 204 is conducted with consideration of other information 206, 216, 214, 210, 208, called non-patent information. The process of assigning

patents to groups is an example of processing patent information with non-patent information. This is the case, because groups are often created according to non-patent considerations. Accordingly, any subsequent processing of the patents in a group involve, by definition, non-patent considerations.

For example, the customer may create groups to represent its products. In this case, the groups are created according to the customer's production information. In another example, the customer may create groups to represent persons of interest. In this case, the groups are created according to HR (human resources) information. In another example, the customer may create groups to represent its competitors. In this case, the groups are created according to business information or practices. In another example, the customer may create groups based on its future products or feature requirement. In this case, the groups are created according to its R&D information.

All of these groups are created based on or in consideration of non-patent information, not patent information. Accordingly, any subsequent group processing of the patents contained in any of these groups represents, by definition, processing of the patent information 204 with consideration of, or in conjunction with, or based on non-patent information 206, 216, 214, 210, 208. This is the case, even if such subsequent group processing involves only, for example, patent bibliographic information (i.e., patent information), such as group processing based on patent issue dates or group processing based on patent references, since the groups being processed were created based on or in consideration of non-patent information, including non-patent information 206, 216, 214, 210, 208.

A group may also contain non-patent documents. In fact, a group may contain only non-patent documents. Accordingly, a group is more generally defined as a collection of documents (such as patent documents only, non-patent documents only, or a combination of patent and non-patent documents). The

-25-

documents in a group typically follow a common theme or characteristic (although this is not a mandatory requirement of groups). Referring to FIG. 2, the invention processes document information 204 alone, or in conjunction with other information 206, 216, 214, 210, 208 (which may or may not be related to the documents). Accordingly, the processing performed by the present invention is more generally described as being document-centric and group-oriented.

Components of the Invention

FIG. 3 is a block diagram of a system 302 according to an embodiment of the invention. The system 302 includes a plurality of databases 316 that store patent information and other information, such as R&D (research and development) information, financial information, licensing information, manufacturing information, HR (human resources) information, and any other information that may be pertinent to the analysis of the patent information. The terms "database" and "table" are used synonymously herein.

An enterprise server 314 accesses and processes the information in the databases 316. In particular, the enterprise server 314 includes modules that are capable of automatically accessing and processing the information in the databases 316 in a patent-centric (or document-centric) and group-oriented manner. These modules are also capable of automatically accessing and processing the information in the databases on a patent by patent basis ("one patent at a time"). Such processing includes, but is not limited to, reporting, analyzing, and planning.

The enterprise server 314 may be a single physical server, or may be a hierarchy of multiple servers 502, 504, 506, 508. An example of this multiple server embodiment is illustrated in FIG. 5. A given client 304, 306 may also connect to one or multiple servers. As is well known, in a client/server

-26-

environment, some work is done by the client, some work is done by the server, and data flows between the client and the server.

The system 302 preferably includes two types of clients, network clients 306 and web clients 304. These clients 304, 306, pursuant to instructions from human operators or users (not shown), interact with the enterprise server 314 to access and process the information in the databases 316. For example, the clients 304, 306 may request that the enterprise server 314 retrieve certain information, or automatically analyze certain information. The enterprise server 314 performs the requested tasks, and sends the results to the requesting clients 304, 306. The clients 304, 306 present these results to their respective operators, and enable the operators to process the results.

Clients 304, 306 may also perform additional processing of data, such as creating a visualization of the data obtained from the enterprise server 314.

Generally speaking, the network clients 306 preferably communicate with the enterprise server 314 using the enterprise server 314's natural language, which is called the enterprise server API (described in detail below). Accordingly, the network clients 306 communicate directly with the enterprise server 314 via a communication network 312, which is preferably a network that uses the well known HTTP (hypertext transport) protocol. Other protocols could alternatively be used. This network 312 may be of any size, such as (but not limited to) a local area network or a wide area network (it can even be a global network).

The web clients 304 do not preferably utilize the enterprise server 314's natural language. Accordingly, the web clients 304 communicate with the enterprise server 314 via a web server 310, which translates between the language of the web clients 304 and the language of the enterprise server 314. This translation is described below.

In an embodiment of the present invention, the components of the present invention shown in FIG. 3 are implemented using well known

-27-

computers, such as a computer 1102 shown in FIG. 11. The computer 1102 can be any commercially available and well known computer capable of performing the functions described herein, such as computers available from International Business Machines, Apple, Silicon Graphics Inc., Sun, HP, Dell,
5 Compaq, Digital, Cray, etc.

The computer 1102 includes one or more processors (also called central processing units, or CPUs), such as a processor 1106. The processor 1106 is connected to a communication bus 1104. The computer 1102 also includes a main or primary memory 1108, preferably random access memory (RAM).
10 The primary memory 1108 has stored therein control logic 1110 (computer software), and data 1112.

The computer 1102 also includes one or more secondary storage devices 1114. The secondary storage devices 1114 include, for example, a hard disk drive 1116 and/or a removable storage device or drive 1118. The removable
15 storage drive 1118 represents a floppy disk drive, a magnetic tape drive, a compact disk drive, an optical storage device, tape backup, ZIP drive, JAZZ drive, etc.

The removable storage drive 1118 interacts with a removable storage unit 1120. As will be appreciated, the removable storage unit 1120 includes a
20 computer usable or readable storage medium having stored therein computer software (control logic) and/or data. The removable storage drive 1118 reads from and/or writes to the removable storage unit 1120 in a well known manner.

Removable storage unit 1120, also called a program storage device or a computer program product, represents a floppy disk, magnetic tape, compact
25 disk, optical storage disk, ZIP disk, JAZZ disk/tape, or any other computer data storage device. Program storage devices or computer program products also include any device in which computer programs can be stored, such as hard drives.

In an embodiment, the present invention is directed to computer program products or program storage devices having software that enables the computer 1102 to perform any combination of the functions described herein.

5 Computer programs (also called computer control logic) are stored in main memory 1108 and/or the secondary storage devices 1114. Such computer programs, when executed, enable the computer 1102 to perform the functions of the present invention as discussed herein. In particular, the computer programs, when executed, enable the processor 1106 to perform the functions of the present invention. Accordingly, such computer programs represent
10 controllers of the computer 1102.

The modules of the invention discussed herein, such as the grouping module 412, the analysis modules 416, etc., preferably represent software executing in the computer 1102.

15 The computer 1102 also includes a display unit 1122, such as a computer monitor, and one or more input devices 1124, such as a keyboard, a mouse, other pointing devices (such as a light pen and trackball), etc.

The computer 1102 further includes a communication or network interface 1126. The network interface 1126 enables the computer 1102 to communicate over communication networks, such as networks 308 and 312,
20 which preferably use the well known HTTP communication protocol.

The components of the invention (shown in FIG. 3) are described in greater detail below. It should be understood that any specific software, hardware, or operating system implementations described herein are provided for purposes of illustration, and not limitation. The invention can work with
25 software, hardware, and operating system implementations other than those described herein. Any software, hardware, and operating system implementations suitable for performing the functions described herein can be used.

Customer Corporate Entity

Preferably, the system 302 is adapted for use by a particular customer. Typically, the customer is a corporate entity. Accordingly, the customer is also called herein the customer corporate entity.

5 It should be understood, however, that the customer can be any organization or individual, such as an academic institution, a research organization, a non-profit or for-profit organization, or any person. Generally, the customer is any entity having an interest in patents.

10 The customer is an entity (such as a company) that has arranged to have use of the system 302 (by purchasing, leasing, or renting the system 302, for example).

15 The databases 316 and data contained therein are specific to the customer. For example, the databases 316 may contain information on the patents that the customer owns and/or licensees, and information on the patents that the customer's competitors owns and/or licenses. Also, the databases 316 may contain the customer's and the customer's competitors' R&D information, financial information, licensing information, manufacturing information, and HR information.

20 Also, the methodology functions supported by the enterprise server 314 may be specialized or augmented to meet the needs of the customer.

25 Implementation and use of the present invention may involve a number of persons associated with the customer corporate entity, such as employees, consultants, associates, and persons retained by the customer, such as attorneys. When interacting with the invention, these people are called operators or users. Table 1 lists some of such persons and their respective responsibilities according to an embodiment of the invention. These persons may be involved in all aspects of the invention for the customer, or may be involved in only some phases of the invention for the customer, such as the extract and load of

-30-

the databases 316. It should be noted that the set up and use of the invention may also involve other people with different knowledge, skills, and/or abilities.

In the discussion contained herein, reference is often made to a user or an operator associated with the customer. It should be understood that the terms "user" and "operator" are synonymous, and refer to one or more persons from Table 1.

Table 1

Role / Function	Responsibilities
Executive, PL, or Division Managers	Ensure strategy meets short and long term business goals and plans
Intellectual Property (IP) Attorneys	Analysis of patents as related to mapping, licensing, infringement, non-renewal, cross-licensing etc.
Technical Personnel	Analysis of patents and how they relate to given product functions and features. Also domain R&D experts as needed for specific competitive technology assessment
MIS personnel	Help in data extraction from operational systems
Marketing personnel	Product Strategy, Features, Target Markets, Competitive Analysis
Business and Financial Analysts	Economic implications, profit, loss, tax, market share, etc.

Databases

FIG. 6 illustrates the databases 316. According to the present invention, the databases 316 store document information (that includes patent information) and information pertinent to the analysis of the document information.

FIG. 6 illustrates a particular embodiment of the databases 316, and also illustrates a particular embodiment of the types of tables that the databases 316 contain, and the attributes in the tables. It should be understood, however, that the invention is not limited to the particular database embodiment of FIG. 6.

5 Instead, the invention is adapted and intended to cover other database structures and organizations that are capable of storing document information and information pertinent to the analysis of the document information. The particular information that is stored in the databases is implementation dependent and varies based on a number of factors, including the type of

10 analysis that is desired, the specific needs of the customer, the type and content of the information that the customer maintains, etc.

The databases 316 of FIG. 6 are collectively called methodology databases, and the data within them are called methodology data, because they support the functions and features, or the methodology, of the present invention. These functions and features (generically called methodology

15 functions and features) are described in sections below.

Many of the databases 316, such as the BOM databases 626, the inventor databases 628, and corporate entity databases 630, the financial databases 638, the person databases 632, and the employee databases 634, are initially loaded using information provided by the customer. Such information includes R&D (research and development) information, financial information, licensing information, manufacturing information, HR (human resources) information, and any other information that may be pertinent to the analysis of

20 the customer's patents and other relevant documents. After initial loading, these databases 316 are updated as necessary to reflect changes in the customer's information.

Other information, such as information for the patent bibliographic databases 604 and the patent database 614, may be loaded using information provided by a third party provider, such as a third party provider that

-32-

specializes in the provision of patent information in electronic form. One such third party provider is SmartPatents Inc. (SPI) of Mountain View, CA. The patent bibliographic databases 604 may be periodically updated through a subscription service from such third party providers. Similarly, the patent database 614 may be augmented through as-needed orders to the third party providers. It should be understood that the present invention works equally well with data provided by any party as long as the data's format matches the formats of the patent bibliographic databases 604 and the patent database 614.

The databases 316 are described in greater detail below.

Document Databases

The document databases 612 preferably include electronic representations of documents of interest to the customer. The document databases 612 represent the customer's repository of documents, and are thus also called the customer's document repository. (The "repository" could alternatively represent all documents represented in the databases 316, whether represented in the document databases 612 or the bibliographic databases 602.)

For example, the patent database 614 includes electronic representations of U.S. and foreign patents of interest to the customer. These patents may be patents owned and/or licensed by the customer, patents owned and/or licensed by competitors of the customer, patents that the customer is considering acquiring, patents that, for whatever reason, the customer is studying, etc. The patent database 614 represents the customer's repository of patents, and is thus also called (in some embodiments) the customer's patent repository.

The patent database 614 preferably has stored therein an image file and a text file for each patent represented in the patent database 614, where the image file and the text file are representations of the patent. Details of an embodiment of the image file and the text file are described in U.S. Patent No.

5,623,681 and U.S. Patent No. 5,623,679, which are both incorporated herein by reference in their entireties.

5 The document databases 612 also include electronic representations of other documents of interest to the customer, such as depositions, pleadings, and prior art references. These documents are respectively stored in a deposition database 618, a pleadings database 616 (generally, pleadings are papers filed with a court), and a prior art database 620. Text and/or image representations of these documents may be stored. These documents may be pertinent to a patent litigation that the customer is involved in.

10 The documents in the document databases 612 may be text, images, graphics, audio, video, multimedia, and/or any other information representation that can be stored in electronic form.

15 It should be understood that the document databases 612 of FIG. 6 are shown for purposes of illustration, and not limitation. As mentioned above, the document databases 612 store electronic representations of documents that are of interest to the customer. Accordingly, the types of document databases 612 and the contents of the document databases 612 are, by definition, customer and implementation specific.

Document Bibliographic Databases

20 The document bibliographic databases 602 store information about documents (as opposed to the documents themselves). More particularly, the document bibliographic databases 602 store bibliographic information about documents.

Patent Bibliographic Databases

The patent bibliographic databases 604 store bibliographic data about U.S. and non-U.S. patents. Such patent bibliographic data includes, but is not limited to, the information on the front page of patents, such as: the patent number, the issue date, the inventors, the title, the assignee, the serial number, the filing date, the U.S. and international classifications, the fields of search, the references cited, the primary examiner, the assistant examiner, the attorney, the agent, the law firm, priority information, related application information, the number of claims, the number of drawing pages, the patent term, the expiration date, etc. The patent bibliographic databases 604 can also include one or more user defined fields that can store large amounts of data, such as 32 Kbytes or more of data.

Operators can extend the bibliographic databases 602 in patent-centric ways. For example, a "current licensee" field can be added to the patent bibliographic databases 604. This could be accomplished, for example, by defining one of the user defined fields to be a current licensee field.

In an embodiment of the invention, the patent bibliographic databases 604 store bibliographic information on all U.S. patents. In other embodiments of the invention, the patent bibliographic databases 604 store patent bibliographic information on a subset of all U.S. patents, such as all U.S. patents that are available in electronic form from the U.S. Patent Office, or all U.S. patents that issued after a certain date.

Generally, there is not a one-to-one relationship between the patents in the patent database 614, and the patents represented in the patent bibliographic databases 604. That is, the patent database 614 does not generally include a copy of each patent represented in the patent bibliographic databases 604. Instead, the patent database 614 includes only those patents that are of interest to the customer. In contrast, the patent bibliographic databases 604 store

-35-

bibliographic information on all U.S. patents and/or foreign patents (or, alternatively, all U.S. patents that issued after a certain date, and/or a subset of foreign patents). Of course, if the customer has an interest in all U.S. patents, such that electronic copies of all U.S. patents are stored in the patent database 614, then there would be a one-to-one relationship between the patents in the patent database 614, and the patents represented in the patent bibliographic databases 604.

Other Document Bibliographic Databases

The document bibliographic databases 602 include store bibliographic information on other types of documents that are of interest to the customer. For example, if the customer is interested in depositions, pleadings, or prior art references, then the document bibliographic databases 602 would store bibliographic information on depositions, pleadings, or prior art references in deposition bibliographic databases 606, pleadings bibliographic databases 608, and prior art bibliographic databases 610, respectively.

The bibliographic information may include the parties or persons involved, the date of creation, the date of modification, the subject, the number of pages, the number of figures, etc. Such bibliographic information may be generated manually, and/or may be generated automatically during the generation of the source document. For example, word processing tools often automatically generate bibliographic information about a document as the document is being created. Such information may include the creator, the typist, the date of creation, the date of modification, the subject, the title, the type of document, the storage format, etc. This automatically-created bibliographic information could be loaded into the document bibliographic databases 602.

Notes Database

The present invention supports annotation of the documents in the document databases 612. More particularly, the present invention allows users to create and link annotations (also called notes) to any portions of the documents in the document databases 612. Such annotations can include text, graphics, images, video, audio, and/or any other information representation that can be stored in electronic form.

The present invention also allows various information to be stored with annotations, such as the date of creation, the creator, the date of modification, a note title and/or subject, access rights, etc.

The annotations, linkage information (i.e., information that specifies the link between a note and a portion of a document), and information related to the annotations and/or the linkage information (such as the position of the linked portion in the document, the date of creation, the creator, the date of modification, a note title and/or subject, access rights, etc.) are stored in the notes databases 640. Embodiments of the notes databases 640 are described in U.S. Patent No. 5,623,679 and U.S. Patent No. 5,623,681, incorporated by reference herein, and in pending U.S. application Ser. No. 08/590,082, which is herein incorporated by reference in its entirety.

Groups Databases

Information on groups is stored in the group databases 621. Generally, a group is a data structure that includes any number of documents that typically follow a common theme or characteristic (although this is not a mandatory requirement of groups). More particularly, a group is a data structure that includes any number of patents that typically follow a common theme or

characteristic (although, again, this is not a mandatory requirement of groups).

Groups are document-centric, or in many cases, patent-centric.

There are two classes of groups: predefined groups (also called system defined groups) and user-defined groups (also called arbitrary groups).

5 However, the invention also supports other types of groups. For example, the invention supports temporary groups. A temporary group is automatically created by the invention in the course of processing a command. One application of temporary groups involves search operations. Specifically, when conducting a search for documents, a new temporary group is created,
10 and the search results are stored in the temporary group. The invention permits operators to convert temporary groups to predefined groups or user-defined groups.

 Patents (and/or documents) in predefined groups follow a predefined theme or characteristic. Database tables, fields, and attributes of a predefined
15 group are specific to the predefined theme/characteristic of the predefined group. Accordingly, different predefined groups have different database tables, different database fields, and different database attributes. Information on predefined groups is stored in the predefined or system defined group databases
20 622.

 Patents (and/or documents) in user-defined groups may or may not follow a common theme or characteristic. Any theme or characteristic that they do follow is defined by the user. Accordingly, user-defined groups are also called arbitrary groups.

 All user-defined groups have the same, generic database tables, fields,
25 and attributes. However, users may elect to use these database tables, fields and attributes differently for different user-defined groups. Information on user-defined groups is stored in the user-defined group databases 624.

 Predefined groups can be more powerful than user-defined groups for at least two reasons. First, the databases associated with a predefined group

store information that is specific to the predefined characteristics of the predefined group. As a result, more useful and specific information can be stored in predefined groups. Second, since the data attributes and characteristics of predefined groups are known in advance, specific functions
5 can be generated in advance to automatically process the information associated with predefined groups. As a result, the information associated with predefined groups can be automatically processed in powerful and diverse ways that are useful given the attributes and characteristics of the predefined groups.

The tables and attributes of predefined groups are typically not
10 applicable to other types of groups. In contrast, the tables and attributes of user-defined groups are generic, and are applicable to all groups. Thus, user-defined groups are more flexible than predefined groups.

Accordingly, in practice, a user-defined group is used by a customer until its attributes, characteristics, and functions are well defined. Once they
15 are well defined, a new predefined group is created to replace the user-defined group. This new predefined group is designed to encompass and take advantage of the specific attributes, characteristics, and functions of the group. In other words, this new predefined group is designed to encompass and take advantage of the well defined structure of the group. Then, analysis and
20 reporting modules are created which automatically analyze and report on the data in the new predefined group. It is possible to create such analysis and reporting modules specific to the new predefined group because of the well defined structure of the new predefined group. The new predefined groups and their reporting and analysis modules can then be distributed (i.e., its databases
25 and functional modules can then be distributed) to interested customers of the invention.

The scope of the present invention includes the creation of new predefined groups and their reporting and analysis functions in the manner described above. The scope of the present invention also includes such new

predefined groups and their reporting and analysis functions. The structure and operation of such new predefined groups and their reporting and analysis functions are implementation dependent, but would be apparent to persons skilled in the relevant art(s) based on the discussion contained herein.

5 In the present invention, groups are structured. Specifically, groups are organized into a directed, acyclic graph, where a group can have multiple children groups and multiple parent groups.

10 The system of the invention discourages or prevents non-sensical organizations of groups. Such non-sensical organizations of groups is at least partially discouraged or prevented by the automatic functions performed by the invention. For example, the system discourages or prevents making a corporate entity group a child of a BOM group, since running an analysis report on all of the subassemblies of the BOM group would yield questionable or undefined results since a corporate entity does not have subassemblies. In an embodiment
15 of the invention, such non-sensical organization of groups is prevented by computer programming.

Also, when a specialized (predefined) group is created to perform specialized analysis functions, new restrictions regarding the rules that govern the inter-relationships between groups are also created. The rules manifest
20 themselves in the database schema. The database schema of the invention prevents the creation of non-sensical group relationships.

Predefined Groups Databases

25 Various predefined groups are described below. It should be understood that the following represents examples of predefined groups supported by the invention. The invention is adapted and intended to include other predefined groups. As described above, predefined groups are often created from user-defined groups once the attributes, characteristics, and functions of the user-

defined groups are well defined. The invention is adapted and intended to include these types of predefined groups. Accordingly, the following is provided for purposes of illustration, and not limitation.

Bill of Materials (BOM) Databases

5 A BOM (bill of materials) group is a group that contains patents (and perhaps other documents) that map to a product, or that map to parts of a product. More particularly, a BOM group is a group that contains patents that map to an assembly, a subassembly, or a part, where an assembly is composed of one or more subassemblies, and a subassembly is composed of other
10 subassemblies and/or parts.

 The phrase "a patent maps to a product" means that the patent includes claims that appear to read on the product or process of making and/or using the product, and/or includes claims that are related to or relevant to the product or process of making and/or using the product, and/or that the patent discloses
15 subject matter than encompasses the product or process of making and/or using the product, and/or that the patent discloses subject matter than is related to or relevant to the product or process of making and/or using the product.

 Information on BOM groups is stored in the BOM databases 626. BOM groups and the BOM databases 626 are discussed in greater detail in sections
20 below.

Corporate Entity Databases

A corporate entity group is a group that contains patents (or other documents) that are owned, licensed, or otherwise of interest to a corporate entity. Information on corporate entity groups is stored in corporate entity
5 databases 630. The corporate entity databases 630 can include information on any number of corporate entity groups. Such corporate entity groups can correspond to any corporate entities that are of interest to the customer, such as the customer itself, affiliates of the customer, competitors of the customer, etc. Corporate entity groups and the corporate entity databases 630 are
10 discussed in greater detail in sections below.

Inventor Databases (and Employees and Person Databases)

An inventor group is a group that contains patents each of which name as inventor a particular person. Information on inventor groups is stored in
15 inventor databases 628. The inventor databases 628 are supported by person databases 632, which include information on people of interest to the customer (people who play a role in the processing of the invention, such as an inventor or employee), and employee databases 634, which include information on employees of interest to the customer. Inventor groups, the inventor databases
20 628, the employee databases 634, and the person databases 632 are discussed in greater detail in sections below.

User-Defined Group Databases

A user-defined group is a data structure that contains documents that follow some user-defined theme or characteristic. Information on user-defined groups is stored in the user-defined group databases 624.

5 These user-defined group databases 624 are common to all user-defined groups. In particular, the attributes in these user-defined group databases 624 are the same for all user-defined groups. However, the customer can choose to utilize these attributes differently for each user-defined group. For example, the customer may choose to store different types of data in these attributes for
10 different user-defined groups. User-defined groups and the user-defined group databases 624 are discussed in greater detail in sections below.

Financial Databases

 The financial databases 638 store financial information pertaining to the customer's business. The financial databases 638 may also include financial
15 information on competitors' businesses (to the extent that such information is publicly known, or can be determined or estimated based on publicly known information or business practices). Such financial information may include money spent on R&D on a product line basis, gross and net revenue on a product line basis, patent licensing revenue, patent acquisition costs, etc. The invention
20 correlates and analyzes the information in the financial databases 638 with patent information to determine, among other things, the financial impact of patents on the customer's and competitors' respective businesses. The financial databases 638 are discussed in greater detail in sections below.

Security Database

The present invention includes multileveled security features for limiting access to data stored in the databases 316. Security is defined herein as privilege levels associated with operators and data objects, and a security methodology for applying the privilege levels so as to restrict access to the data objects to operators having the appropriate privilege levels.

The invention is capable of supporting security for all data items, including security for notes (stored in the notes databases 640), groups (stored in the group databases 621), financial information (stored in the financial databases 638), personal information (stored in the person databases 632 and the employee databases 634), and documents (stored in the document databases 612 and the document bibliographic databases 602). Information for implementing these security features is stored in the security databases 636, which are discussed in greater detail in sections below.

Enterprise Server

The enterprise server 314 is preferably implemented as one or more computers (such as the computer 1102 shown in FIG. 11) each having at least 128 MBytes of main memory 1108 and running Microsoft Windows NT. The enterprise server 314 could, alternatively, be implemented using other memory configurations, and other operating systems, such as (but not limited to) UNIX, Windows 95, MS-DOS, the Apple Operating System, etc. Accordingly, the specific hardware and software implementations discussed herein are provided for purposes of illustration, not limitation (this applies to all specific hardware and software implementations discussed herein, both for the enterprise server 314 and for other components of the invention). The invention can utilize any hardware,

software, and operating system capable of performing the functions described herein.

The enterprise server 314 can be a single computer, or a hierarchy of multiple computers (FIG. 5). Logically, however, the enterprise server 314 is preferably a single computer.

FIG. 4 is a logical block diagram of the enterprise server 314. The enterprise server 314 has a number of modules (collectively called the enterprise server modules). Note that a number of the modules interact with the databases 316. A SQL server 426 (such as the Microsoft SQL Server) and/or other well known database servers 428 interact directly with the databases 316. The enterprise server modules interact with these servers 426 and 428 and the databases 316 via a database interface module 420, which preferably represents an ODBC (object database connectivity) layer.

The Network transport layer or interface 401 is used to receive command request objects from the client 304, 306 based on a specific network protocol, preferably HTTP. On the enterprise server 314 these network command objects are reconstructed from a stream of bits received from the client 304, 306. Once the command objects have been reconstructed the specific operations (described herein) defined in this object are performed by the appropriate enterprise server modules. The command objects represent enterprise server API commands, discussed below.

According to an embodiment of the invention, command objects include autonomous intelligent agents that perform appropriate operations at the enterprise server 314 on behalf of the operator (i.e., the client 304, 306). In this embodiment, the command objects sent to the enterprise server 314 represent computer programs that are executed in the enterprise server 314. These executing computer programs preferably represent threads each having an address space. These computer programs, when executing in the enterprise server 314,

perform the functions discussed herein, such as patent mapping, patent aging, inventor count, inventor information, financial functions, etc.

The enterprise server 314 is a highly secure business decision system. The specific operations in each command object are checked against the security information maintained about each user in the system. This is logically done through a comprehensive security layer or module 402. (The specific implementation of security requires the interaction with ODBC 420, as all security information is stored in the databases 316). Alternatively, the security module 402 could logically be shown as being under the server configuration module 404 and the command dispatch module 406.

As described elsewhere herein, the document storage and retrieval module 408 is part of a Virtual Patent System 11304 (FIG. 113) that presents a consistent, unified view of an arbitrary number of patent and patent-related documents.

The Searching subsystem or module 410 provides for patent searching using a search language (syntax) described below, an extensible language for searching patent and other patent-related documents. The search layer 410 also encapsulates the specific search engine 424 used in the implementation of the system, which can and will vary based on available search technologies.

The other layers shown in FIG. 4 work together to form the heart of the business decision system of the present invention. The Groups layer or grouping module 412 is responsible for managing all groups created by a user in support of patent analysis. The Notes layer or module 414 is responsible for managing all forms of annotations made by the user. The Analysis Queries layer or analysis modules 416 perform analysis queries in support of specific requests made by various modules in the decision support system. Finally, the server administration layer or module 418 provides services to manage the configuration of the enterprise server 314, such as adding or changing the security permissions associated with a specific user.

-46-

Each of these layers provides a mechanism to further decouple the operation of the enterprise server 314 from the specific implementation of the databases 316. Each of these layers also interact with ODBC (Open Database Connectivity) 420, a Microsoft defined industry standard mechanism for manipulating relational databases (other software for interacting with and manipulating databases could alternatively be used). ODBC 420 provides a final layer of decoupling and enables the enterprise server 314 to transparently connect to different relational databases 316.

The enterprise server modules are further described below.

Document Storage and Retrieval Module

The document storage and retrieval module 408 in the enterprise server 314 stores and retrieves documents from the document databases 612. Preferably, especially with respect to patent documents, the document storage and retrieval module 408 stores and retrieves text files and image files representative of documents in the document databases 612. The document storage and retrieval module 408 performs such data storage and retrieval operations pursuant to commands that conform to the enterprise server API, described below.

The document storage and retrieval module 408 preferably interacts directly with the operating system 422 of the enterprise server 314, where such direct interaction primarily pertains to data retrieval and storage.

As just noted, the document storage and retrieval module 408 operates to access data in the document databases 612, such as the customer's repository of patents represented by the patent database 614. Preferably, the patent database 614 stores electronic representations of all patents which are of interest to the customer. Additional electronic patents can be added to the patent database 614 at any time as the customer's interests change. The patent database 614 is capable

of storing electronic representations of all U.S. patents, or any subset of all U.S. patents, and of any number of foreign patents as required by the customer's needs and interests. Accordingly, the document storage and retrieval module 408, in combination with the patent database 614 and the patent bibliographic databases 604, provide the customer with the ability to quickly, efficiently, and effectively access, display, and process any patent of interest. Accordingly, from the perspective of the client, the document storage and retrieval module 408, in combination with the patent database 614 and the patent bibliographic databases 604, represent a virtual patent system. FIG. 113 graphically depicts this virtual patent system 11304.

The client document storage and retrieval module 708 in the clients 304, 306 (FIG. 7) displays the text and images received from the document storage and retrieval module 408 in the enterprise server 314. As shown in FIG. 112, the client document storage and retrieval module 708 is capable of simultaneously displaying the text of a document in a first window 11202, and the image of a document in a second window 11204.

The client document storage and retrieval module 708 has features and functions for enabling a user to manipulate and otherwise process the displayed data. For example, the client document storage and retrieval module 708 includes text searching features, powerful text and image navigation features, text processing features, image processing features (as represented by image toolbox 11206 shown in FIG. 112), document organization features, word list features, sophisticated text and image display features, text and image highlighting features, document importation and exportation features, case or group copying features, and print features.

The document storage and retrieval module 408 in the enterprise server 314 and the client document storage and retrieval module 708 in the clients 304, 306 are collectively further described in U.S. Patent No. 5,623,679 and U.S. Patent No. 5,623,681, incorporated by reference herein, and in pending

U.S. application Ser. No. 08/341,129, which is herein incorporated by reference in its entirety.

Notes Module

5 The notes module 414 manages and interacts with the notes databases 640. The notes module 414 processes enterprise server API commands (described below) to: create new notes, update existing notes, add notes to a document, remove notes from a document, and retrieve all notes associated with a document.

10 The client notes module 714 enables a user to view and manipulate notes. FIG. 111 is a screen shot displayed by the client 304, 306 on the client monitor 1122. Text of a patent is displayed in a first window 11104. The client notes module 714 displays upon command the notes that are linked to portions of this patent in a notes window 11108.

15 The client notes module 714 receives from the user commands to, for example, edit note contents, create new notes, link new or existing notes to portions of documents, modify notes, and delete notes. The client notes module 714 modifies the display of the notes window 11108 as necessary to reflect these user commands. The client notes module 714 also generates enterprise server API commands corresponding to these user commands, and
20 forwards these enterprise server API commands to the enterprise server 314 for processing by the notes module 414 in the enterprise server 314.

25 Notes may have attributes, such as (but not limited to) the person who created the notes (relevant for security purposes), the date the note was created, the data format(s) of data stored in the note (text, image, graphics, video, audio, spreadsheet, database, etc.), the note title, the note subject, whether the note contains information that would be considered to be Attorney/Client privileged or confidential, and the date the note was last modified.

According to an embodiment of the invention, notes are hierarchically organized. That is, a given note may be a child note of any number of parent notes, and may have any number of child notes. This, of course, is in addition to the linkage of notes to portions of documents. This hierarchical organization may be implemented by having in the note databases 640 a note_note_xref table, that would be similar to the group_group_xref table 1229. The note_note_xref table would have a parent note attribute storing the note ID of the parent note, and a child note attribute storing the note ID of the child note. There would be a record in the note_note_xref table for each parent note/child note relationship in the note hierarchies. It is noted that this note hierarchy provides a structure, organization, and hierarchy to the documents linked to the notes.

The notes module 414 in the enterprise server 314 and the client notes module 714 in the client 304, 306 are collectively further described in U.S. Patent No. 5,623,679 and U.S. Patent No. 5,623,681, incorporated by reference herein, and in pending U.S. application Ser. No. 08/590,082, incorporated by reference herein.

Searching Module

The searching module 410 in the enterprise server 314 interacts with a search engine 424 to conduct searches through the data in the databases 316 pursuant to search requests from the clients 304, 306. The search engine 424 is any commercial and well known search engine. Preferably, the search engine 424 is implemented as the Fulcrum search engine available from Fulcrum Technologies, Inc., Ottawa, Canada. Other commercial search engines could also be used, including (but not limited to) those from Verity Incorporated, Sunnyvale, CA, Open Text of Canada, and others.

Preferably, the data in the databases 316 is indexed to facilitate and enhance searching by the search engine 424. For example, each field in each table of the patent bibliographic databases 604 is preferably indexed and searchable. Also, the documents (including the text files and possibly the image files) in the document databases 612 are preferably indexed and searchable. Any well known indexing procedure can be used to index the data in the databases 604. According to an embodiment of the invention, indexing and searching are performed as described in pending U.S. patent application Ser. No. 08/422,528, which is incorporated herein by reference in its entirety. Searching for documents is performed by searching through these indexes. The index tables are preferably stored in the searching module 410, in the searching engine 424, and/or in the databases 316.

An embodiment of the invention permits operator-defined indexing of data. In this embodiment, an operator can define what data in the databases 316 is to be indexed. For example, an operator can specify that only patents having as assignee "IBM" should be indexed. Or, the operator can specify that only the documents in a given group should be indexed. Such operator-defined indexing enhances searching performance, because the index that is searched is smaller and more targeted.

The searching module 410 receives enterprise server API commands from the clients 304, 306. The searching module 410 processes these enterprise server API commands and, as a result, causes the search engine 424 to perform at least the following functions: conduct a search to identify documents that satisfy a client-supplied search parameter (for example, to identify documents that contain instances of key words), retrieve and return the search results of a previously executed search, and retrieve and return search hit information for a particular document so that search term highlighting can be performed on the document.

According to the present invention, the documents identified by a search can be easily added to a new group or an existing group by invoking appropriate enterprise server commands, such as the ReqAddDocListToGroup command or the ReqAddPatents command. In the user interface at the client
5 304, 306, the operator implements this function using drag-and-drop techniques.

Preferably, the invention creates a new, temporary group to hold the results of a search. A subsequent search could then be scoped or restricted to the documents in this temporary group. Accordingly, the invention supports
10 iterative searching using groups.

The invention supports many search strategies, including but not limited to keyword, keyword phrase, keyword phrases with boolean, thesaurus, concept searching, object searching, and graphical searching based on likeness of words/images.

The client searching module 710 in the clients 304, 306 receives search
15 commands from the user. The client search module 710 converts these search commands to corresponding enterprise server API commands, if necessary, and transfers these enterprise server API commands to the enterprise server 314. The client searching module 710 receives from the enterprise server 314 search
20 results. The client searching module 710 displays these search results and enables the user to manipulate and process the search results (such as by enabling the user to add the documents identified by a search to a new or existing group -- note that this functionality may also involve the client grouping module 712).

The invention also supports restricting or defining a search according to
25 aspects of the system, such as historical information. Such historical information can include, for example, the results of a prior search. Thus, the scope of a new search can be restricted to the results of a prior search, or the search criteria in a new search can be added to the search criteria in a prior

search. Preferably, the system maintains a search log so that the operator can view and select prior search results and prior search criterions.

In some embodiments, a user's characteristics (i.e., security level) define the groups that the user can search in. In other words, searches are restricted to groups for which the user has access rights.

The operation of the client searching module 710 in a client 304, 306 and the searching module 410 in the enterprise server 314 shall now be described in greater detail with reference to FIG. 9. The client searching module 710 supports a number of user interfaces for enabling the user to enter a search command. One user interface is a field driven graphical user interface (GUI) 902. Examples of field driven GUIs 902 are shown in FIGS. 53 and 57.

Considering first FIG. 53, the client searching module 710 displays the searching window 5302 on the client display monitor 1122. The searching window 5302 includes a Scope of Search field 5304 through which the user can select a scope of search. The user presses a down-arrow button 5306 to obtain a list of possible search scopes. This list may include, for example, all U.S. patents, all foreign patents, both U.S. and foreign patents, all patents in one or more selected groups, the patents in the customer's repository, etc. Searches can also be restricted to portions of documents, such as the claim section in patents. In the example of FIG. 53, the user has defined the search scope as being all U.S. patents.

The fields in the searching window 5302 allow the user to specify a search of patent bibliographic information, and/or a search of the text of patents. The user can search through patent bibliographic information by entering key terms in the patent number field 5306, the title field 5308, the inventor field 5310, the assignee field 5312, the class field 5314, and/or the date of issue field 5315. The date of issue field 5315 allows the user to specify patents that issued before or after a given date (by filling in fields 5316 and 5318), or that issued between two dates (by filling in fields 5320, 5322, and

5324). It is noted that only some of the attributes of the patent bibliographic databases 604 are shown as being searchable in FIG. 53. In some embodiments, other field driven GUIs (not shown) supported by the invention have search fields corresponding to other attributes of the patent bibliographic databases 604. In these other embodiments, it is possible for the user to search through any of the attributes of the patent bibliographic databases 604.

The user can search through the text of patents by entering search parameters in an abstract field 5326 and/or the full patent text field 5328.

It is noted that not all users may have access to all of the search options described above. For example, some users may be only able to search through the patent bibliographic information. Other users may be only able to search through certain attributes of the patent bibliographic information. Other users may be only able to search through the text of patents. The server configuration module 404, described below, controls the search options and capabilities of each user.

The user can specify the fields to include in the list of search results by appropriately selecting fields 5330. The user can specify a sorting order to display the search results via field 5332. Sorting options include: descending patent numbers, ascending patent numbers, issue date, filing date, serial number, score (the number of search hits), etc.

FIG. 54 illustrates an example screen shot of search results displayed by the client searching module 710 on the client display monitor 1122. By selecting a "get results in a file command" 5406, the user can write the search results to a user-specified file. By selecting a "patents in local repository command" 5408, the user can display a list of the patents from the search results that are stored in the patent database 614 (i.e., whose text and image files are stored in the patent database 614).

By selecting a "patents not in local repository" command 5410, the user can display a list of the patents from the search results that are not represented

in the patent database 614 (i.e., patents for which the user does not own electronic copies of). The report resulting from selecting the patents not in local repository command 5410 can be used by the user to generate a purchase order to obtain electronic copies of the patents of interest from the search results. In some embodiments, electing this option will cause an electronic message to be sent to a third party service provider. The third party service provider would then electronically send electronic copies of the patents to the customer.

If the user selects (by double clicking or other well known GUI operation such as selecting a patent and pressing a return button) a patent from the list shown in FIG. 54, then the text and/or image of the selected patent is displayed on the client display monitor 1122. FIG. 55 depicts the display of text, and FIG. 56 depicts the display of an image. Alternatively, both the text and image can be simultaneously displayed on at least some clients 304, 306 using a display format such as that shown in FIG. 112.

The field driven GUI 5702 of FIG. 57 is similar to that of FIG. 53. Note that the GUI 5702 of FIG. 57 includes a keywords field 5716, which allows the user to search through user-definable fields in the patent bibliographic databases 604. The field driven GUI 5702 of FIG. 57 also allows the user to define the scope of the search via fields 5728. In the example of FIG. 57, the scope of the search can be the full text index (i.e., a search of the patent bibliographic information), only the patents stored in the patent database 614 (i.e., only the patents in the customer's patent repository), only the patents in the current group, or only the current patent. Other embodiments may restrict searching to specific types of documents or specific predefined groups, such as all European patents, all PCT applications, all non-patent documents, documents in BOM groups, etc.

Referring again to FIG. 9, the client searching module 710 generates a query request 908A based on the search criteria that the user entered into the

-55-

field driven GUI 902. Preferably, this query request 908A is in the native query language of the enterprise server 314. In other words, the query request 908A conforms to the enterprise server API.

The enterprise server API commands related to querying include the ReqSearch command. As described further below, this command takes searchParameters as a passed parameter. This passed parameter stores the search parameters for the search. A preferred syntax of the search parameters according to the enterprise server API is described below in Tables 2 and 3.

Table 2

Search string operator	Meaning and Search Behavior	Implementation in Search Engine 424 (when using Fulcrum as the Search Engine 424)
W/n	Search for term expression on left within n characters distance in either direction from term expression on right.	Translate directly to "WITHIN n CHARACTERS OF"
AND	Match only documents that satisfy the term expression on the left and the term expression on the right.	Translate directly to "&"
OR	Match documents that satisfy the term expression on the left or the term expression on the right.	Translate directly to " "

-56-

NOT	Match only documents that do not satisfy the term expression on the right.	Translate directly to "!"
()	Parentheses. Used to group search expressions parts to control their order of evaluation.	Leave as is.

Each of the Operators in Table 2 (including any spaces to its immediate left or right) is considered to be a search syntax delimiter. Each sequence of characters before, after, or between one of these delimiters will be called a search string "element". Each search string element will be enclosed between a pair of apostrophes to translate it for transmission to Fulcrum. The meaning of and translations for the specific characters that can appear in an element are listed below in Table 3.

Table 3

Search string Element	Meaning and search behavior	Implementation in Fulcrum
A-Z a-z	Alphabetic characters. A contiguous sequence of these (including any optional apostrophes) is considered a word for searching. All searching is case insensitive.	Leave as is.

0-9	Numeric characters. A contiguous sequence of these (including any optional commas or periods) is considered a word for searching.	Leave as is.
'	Apostrophe. This character only appears in the index when there is an alphabetic character on either side of it. In this case, you must search for it explicitly. For example, searching for "Adams" will not find "Adam's".	Translate directly to "'"
,	Comma and period. Each of these characters only appears in the index when there is a numeric character on either side of it. In this case, you must search for it explicitly. For example, searching for "4,234.03" will not find "423403".	Leave as is.
*	Wildcard matching zero or more characters in a single word.	Translate directly to "%"
?	Wildcard matching exactly one character in a single word.	Translate directly to "_"
%	Fulcrum's wildcard matching zero or more characters in a single word.	Translate directly to "%\"

-58-

_	Fulcrum's wildcard matching exactly one character in a single word.	Translate directly to "_"
\	Escape character in Fulcrum.	Translate directly to "\\ "
Space	Space character.	Leave as is.
-	Behaves like the "other punctuation" below with the exception that when one or more dashes appear in the middle of a word in a search string, the search engine will search for both the version with all the dashes and the version with none of the dashes.	Leave as is.
!@#\$%^&_- =+[]{};< >"/ `~	All other punctuation. These are treated as invisible word breaks. They are not indexed, but will break words.	Leave as is.

The searching module 410 in the enterprise server 314 receives the query request 908A. A query language syntax analyzer 914 in the searching module 410 checks the query request 908A for any format or syntax errors, such as unbalanced parentheses. The searching module 410 then translates the query request 908A to a new query request in the language of the search engine 424. The new query request is then transferred to the search engine 424 for processing.

The present invention also supports a native language command line GUI 904 for enabling a user to enter a search request. The command line GUI

904 is typically only used by users who are familiar with the enterprise server API. When using the command line GUI 904, the user enters at the command line a query request 908B. This query request 908B must conform to the enterprise server API. This query request 908B is then transferred to the searching module 410 in the enterprise server 314 where it is processed in the manner described above.

The present invention further supports any number of foreign language command line GUIs 906 for enabling the user to enter query requests. The invention provides foreign language command line GUIs 906 to support those users who are familiar with database query languages other than the enterprise server API. Such database query languages are herein called foreign query languages for reference purposes. There are many well known foreign query languages, such as the patent specific query language used by the U.S. Patent Office Web Site which is located at <http://patents.cnidr.org/access/access.html>. The client searching module 710 has a foreign language command line GUI 906 for each foreign query language of interest.

When using a foreign language command line GUI 906, the user enters at the command line a query request 910. The query request 910 is in the foreign query language associated with the foreign language command line GUI 906. The query request 910 is translated to a query request 908C in the enterprise server API by a translator 912 (there is a translator for each foreign query language supported by the invention). This query request 908C is then transferred to the searching module 410 in the enterprise server 314 where it is processed in the manner described above.

The present invention also supports searching of other data objects, such as groups (in the group databases 621) and notes (in the notes databases 640). In fact, the present invention supports searching of all the tables in the databases 316. Preferably, all fields in all tables of the databases 316 are indexed and searchable. In some embodiments, only some of the tables are

indexed and searchable, such as the group databases 621 and the notes 640. GUIs, such as those discussed above, are used to enable operators to define searches of any attributes of these tables.

5 The present invention also supports context and linguistic type searching, and also supports image and object searching. The invention can be used, for example, with data blade search tools, such as those available from Informix.

Automatic Searches Related to Groups

10 The present invention also supports an automated search function related to groups. According to this aspect of the invention, a search is performed of all or part of the document databases 612 and/or the document bibliographic databases 602 to identify documents that satisfy a specified search criteria. The documents identified via this search are added to a specified group.

15 For example, suppose that the customer has a group called XYZ group. This group contains the patents that name XYZ corporation as assignee. Periodically, the invention automatically searches the patent bibliographic databases 604 for any patents that name the XYZ corporation as assignee. Any patents found from this search are automatically added to the XYZ group.

20 The invention supports performing such automatic searches at user defined intervals (such as every month), or at the occurrence of user-specified events, such as whenever the patent bibliographic databases 604 are updated.

25 The invention allows the customer to define such automatic searches. In defining an automatic search, the customer specifies the target databases (what databases to search), the target groups (which groups receive the identified documents), the search criteria, and the frequency or circumstances that the automatic searches take place.

Preferably, the searching module 410 performs the automatic searches.

Searching Algorithm

The searching module 410 processes a search string according to a preferred searching algorithm that is designed to take advantage of the searching and data accessing capabilities of the objects that directly interact with the databases 316. Such objects are herein called database accessing objects because they directly access and interact with the databases 316, and include the search engine(s) 424, the SQL server 426, and other database servers 428.

A flowchart 13902 shown in FIG. 139 represents a searching algorithm performed by the searching module 410 according to a preferred embodiment of the present invention. The searching module 410 performs the steps of flowchart 13902 with respect to a search string that it has received from a requester, such as a client 304, 306, or any other entity that wishes to conduct a search of the databases 316.

The search string includes one or more search string components, also called search string elements, which are preferably in the format shown in Table 3. The search string components/elements are separated by search syntax delimiters (Table 2).

In step 13906, the searching module 410 identifies the search string components in the search string. The searching module 410 preferably performs step 13906 by parsing through the search string. In the course of such parsing, the searching module 410 identifies search string components based on the location of search string delimiters (that is, search string components represent groups of characters that are separated by search string delimiters). For example, consider the following example search string:

(Phrase1 AND Phrase2) OR (Phrase6 AND (Phrase3 OR (Phrase4 AND Phrase5))).

In step 13906, the searching module 410 identifies the following as search string components (parsing the example search string from left to right):

-62-

Phrase1, Phrase2, Phrase6, Phrase3, Phrase4, and Phrase5.

In step 13908, the searching module 410 analyzes each search string component (identified in step 13906) and assigns each search string component to a database accessing object. The searching module 410 in step 13908 assigns
5 a search string component to a database accessing object based on the characteristics of the search string component and the capabilities of the database accessing object. Specifically, the searching module 410 analyzes and identifies the characteristics of the search string component. The searching module 410 then assigns the search string component to a data accessing object whose
10 capabilities are matched to these characteristics (that is, whose capabilities are well suited for processing search string components having those characteristics).

For example, suppose that the search string component being considered represents a text search in a collection of documents. This type of search is best suited to be performed by a search engine, such as search engine 424. Search
15 engine 424 is well suited for performing text searches because the text in the databases 316 is indexed.

As another example, suppose that the search string component represents a search for all patents that are referenced by U.S. Patent No. 1,234,567. This
search string component is best represented as a relational database query.
20 Accordingly, it would be best processed by a relational database engine, such as the SQL server 426 or other relational database servers 428.

After the search string components have been assigned to data accessing objects, the data accessing objects in step 13908 process their assigned search string components. Such processing preferably occurs in parallel. By processing
25 the search string components in parallel, the length of time that it takes to conduct the search is reduced.

In step 13910, the data accessing objects transfer their respective result sets or search results to the searching module 410. These search results represent multiple data streams. The searching module 410 in step 13910 combines these

data streams according to the logical linking operations represented by the search syntax delimiters in the original search string.

For example, FIG. 138 illustrates the manner in which the searching module 410 combines the search results resulting from the example search string presented above. R1 represents the search results generated by processing
5 Phrase1. Similarly, R2, R3, R4, R5, and R6 represent the search results generated by processing Phrase 2, Phrase3, Phrase4, Phrase5, and Phrase6.

The searching module 410 combines R4 and R5 by logically AND'ing R4 and R5. That result is OR'd with R3. The result of that OR operation is AND'd
10 with R6. The result of that AND operation is OR'd with the result of the logical AND'ing operation of R1 and R2. The result of this OR operation represents the final result of processing the original search string.

Preferably, for efficiency purposes, the result sets received from the database accessing objects are ordered according to a common criteria.
15 Preferably, the result sets are ordered according to patent number.

The searching module 410 in step 13910 then returns this final result to the requester. Operation of flowchart 13902 is then complete, as indicated by step 13912.

In the present invention, it is not necessary to store the intermediate search results. For example, it is not necessary to store R1, R2, R3, R4, R5, and R6.
20 Instead, the searching module 410 processes the intermediate search results as they are received. Referring to FIG. 138, the searching module 410 processes R4 and R5 (as they are received) by AND'ing them together. The result of that AND operation is immediately processed with R3 (as it is received) by OR'ing them
25 together. The other search results are processed as they are received in a similar manner. As a result, the searching module 410 does not need to store the intermediate search results R1-R6 for any length of time.

The searching algorithm of the present invention shall now be further described with reference to the processing of the "patents in local repository"

-64-

command 5408 and the "patents not in local repository" command 5410 (FIG. 54). These commands were discussed above.

5 The searching module 410 executes the "patents in local repository" command 5408 by first processing a search string in the manner shown in the flowchart 13902 (FIG. 139). The results of processing this search string are then logically AND'ed with a listing of the patents in the patent repository (that is, in the patent database 614). This logical AND operation yields a list of patents which satisfied the search string, and which are also in the patent database 614.

10 The searching module 410 performs the "patents not in local repository" command 5410 by first processing a search string in a manner discussed above with reference to flowchart 13902 in FIG. 139. The results of processing the search string are then logically NAND'ed with the list of patents in the patent database 614. This NAND operation yields a list of patents which satisfy the search string, but which are not in the patent database 614.

15 *Grouping Module*

20 The grouping module 412 in the enterprise server 314 manages and interacts with the group databases 621. The grouping module 412 receives and processes enterprise server API commands (sent from clients 304, 306) to perform at least the following functions: obtain information on the hierarchy of the groups stored in the group databases 621, make an existing group a child of another group, unlink a child group from one of its parent groups, update group properties, create a new group as a child of an existing group, obtain a list of documents in a group, add documents to a group, and remove documents from a group.

25 The grouping module 412 also supports group import and export functions. Some of these group import functions relating to BOM groups, corporate entity groups, and inventorship groups are described below. The

-65-

grouping module 412 also supports a user-defined group (or generic) group import function. When performing this function, the grouping module 412 receives a properly structured file comprising a plurality of records (or lines of information), where each record specifies a name of a group and a physical level of the group in the group hierarchy. From this file, the grouping module 412 creates a user-defined group hierarchy in the user-defined group databases 624.

An example file for group importation is shown below (other file structures could alternatively be used):

10	1	A
	2	B
	3	C
	2	D
	3	E
15	4	F
	3	G

The first column corresponds to the physical level in the group hierarchy, and the second column corresponds to the group name. The information contained in this example file corresponds to a group hierarchy shown in FIG. 115.

The grouping module 412, upon receipt of this file, creates a record in the group_table 1227 for each of the groups represented in the file (i.e., for groups A, B, C, D, E, F, and G). The physical level information from the file is stored in the group_group_xref table 1229. For example, the group_group_xref table 1229 would have a record that indicates that Group A is the parent of Group B. The group_group_xref table 1229 would also have a record that would indicate that Group G is a child of Group D. (It is noted

-66-

that similar physical level information is also preferably in the customer BOM data 4704 (FIG. 47) with respect to load of the BOM databases 626.)

The client grouping module 712 displays the group hierarchy and the documents in a group, and enables a user to manipulate and process groups.

5 FIG. 58 depicts an example screen shot displayed by the client grouping module 712 on the client monitor 1122. In a first window 5804, the client grouping module 712 displays a graphical representation of the hierarchy of groups stored in the group databases 621. Suppose that the user has selected an ALU group 5808 in this first window 5804. Selection of a group in the first window 5804
10 causes a list of the documents in the selected group to be displayed in a second window 5806. Accordingly, the client grouping module 712 displays the following list of documents: U.S. Patent Nos. 4,333,AAA; 4,788,BBB; 5,278,CCC; 4,760,478 (as should be apparent from this example, many of the patents referred to herein for illustrative purposes are fictional). These
15 documents are in the selected ALU group. Note that the second window 5806 also displays bibliographic information on the listed documents. Preferably, the information listed in the second window 5806 is in a spread sheet format. However, other formats could alternatively be used.

Selecting (by double clicking) a document from the list in the second
20 window 5806 causes the selected document to be displayed. For example, suppose that the user selected U.S. Patent No. 4,760,478 from the list displayed in the second window. This would cause the client 304, 306 to obtain the text and image of this patent from the databases 316 via the enterprise server 314. The client document storage retrieval module 708 would
25 then display the retrieved text and image of the '478 patent at the client monitor 1122 using the format shown in FIG. 112, where the text is displayed in a text window 11202, and the image is displayed in an image window 11204.

The client grouping module 712 receives from the user commands to navigate through the group hierarchy, to edit the group hierarchy, to edit

-67-

groups, to add documents to groups, to delete documents from groups, to delete groups, etc. The client grouping module 712 modifies the display of the window 5802 as necessary to reflect these user commands. The client grouping module 712 also generates enterprise server API commands corresponding to these user commands, and forwards these enterprise server API commands to the enterprise server 314 for processing by the grouping module 412 in the enterprise server 314.

Analysis Modules

The analysis modules 416 are shown in FIG. 10. These analysis modules 416, which are also called methodology modules 416, automatically interact and process data contained in the databases 316 pursuant to user commands. The analysis modules 416 are patent-centric (or document-centric) and group-oriented. The analysis modules 416 are patent-centric because they all involve the processing (including reporting, analysis, and planning) of patent data either with or without consideration of other data, such as production data, HR data, financial data, etc. The analysis modules 416 are group-oriented because they have the capability of processing the patents (or other documents) in one or more groups, and potentially the children of these groups.

It should be understood that the invention is adapted and intended to include a wide and varied range of analysis modules 416. The analysis modules 416 shown in FIG. 10 represent only a sampling of the analysis modules 416 that the invention is adapted and intended to support. The invention can support many other analysis modules 416 because the databases 316 are so rich. The analysis modules 416 can include any other module that performs useful processing (from the point of view of the customer) of the data in the databases 316. Accordingly, the particular collection of analysis modules

-68-

416 shown in FIG. 10 are described herein for the purpose of illustration, and not limitation.

For illustrative purposes, the analysis modules 416 are sometimes described herein as working with particular types of groups, such as BOM groups. However, it should be understood that the analysis modules 416 can work with any types of groups.

The analysis modules 416 are described in detail in sections below.

Security Module

The security module 402 in the enterprise server 314 manages and interacts with the security databases 636, which stores information required to implement the security features of the invention. The security module 402 utilizes the information in the security databases 636 to implement a multilevel security methodology. The security databases 636 and the multilevel security methodology implemented by the security module 402 are described in detail in sections below.

The client security module 702 in the clients 304, 306 enable a user to access and modify the security information in the security databases 636. Typically, access to the client security module 702 and the security databases 636 is limited to users with high security clearances, such as system administrators.

Server Administration Module

The server administration module 418 performs functions related to the enterprise server 314's resources, such as the databases 316 and the indexes to the databases 316. Some of the server administration module 418 functions include performing reindexing operations of the databases 316, when necessary,

importing and exporting large portions of the databases 316, upon request, managing directories, etc. The server administration module 418 is also responsible for establishing user sessions with the enterprise server 314.

5 The client server administration module 718 at the client 304, 306 and/or the server administration module 418 in the enterprise server 314 preferably maintains one or more server configuration files. Information in the server configuration files identifies, for example, the physical location of the databases, the number of possible concurrent users, memory size allocations, etc. The server configuration files can also include log files. The server
10 configuration files would indicate what events are to be logged to the log files, such as whether to track all user actions, track error conditions, etc.

Server Configuration Module

Not all users have access to all of the enterprise server functions. A normal user, for example, may be only able to search documents and document
15 bibliographical information, view documents, and print documents. Accordingly, a normal user would be able to access only the document storage and retrieval module 408 and the searching module 410 of the enterprise server 314.

Also, different users may have different search capabilities. Some users
20 may be able to search only document bibliographic information, while other users may be able to search both document bibliographic information and the text of documents.

A power user may be allowed access to all enterprise server functions. The power user could search documents and document bibliographical
25 information, view documents, print documents, work with groups, work with notes, and invoke analysis functions. Accordingly, the power user would be able to access the document storage and retrieval module 408, the searching

-70-

module 410, the grouping module 412, the notes module 414, and the analysis modules 416 of the enterprise server 314.

A system administrator may be able to set user security levels and perform enterprise server administration functions. Accordingly, a system administrator would have access to at least the security module 402 and the server administration module 418.

The modules loaded on a user's computer preferably depend on the user's security level, and correspond to the modules in the enterprise server 314 to which the user has access. Referring to FIG. 7, only the client document storage and retrieval module 708 and the client searching module 710 are preferably loaded on the computers (i.e., clients 304, 306) of normal users. The client document storage and retrieval module 708, the client searching module 710, the client grouping module 712, the client notes module 714, and the client analysis module 716 are loaded on the computers of power users. The client security module 702 and the client server administration module 718 are preferably loaded on the computers of system administrators.

The system configuration module 404 in the enterprise server 314 keeps track of the functions and modules that each user is permitted to access. Preferably, the server configuration module 404 maintains a database (not shown) having access privilege information for each user (note that this is different than security access information). As shown in FIG. 101, when the enterprise server 314 receives a command from a user, the server configuration module 404 in step 10108 ensures that the user has access to the target module (i.e., the module in the enterprise server 314 that would have responsibility for processing the command) before allowing the command dispatch module 406 to forward the command to the target module for processing. FIG. 101 is further described below.

The functions and enterprise server modules that a user is allowed to access is dependent on a number of factors, such as the user's level of need, the

user's level of expertise, and/or whether or not the user has purchased the modules and/or databases (in some embodiments, it may be necessary for the user to pay a fee to access and obtain the benefits of an enterprise server module, such as the notes module 414, the grouping module 412, and/or the analysis modules 416).

The user's computer platform is also a consideration. In some embodiments, software for some client modules (such as the client grouping module 712, the client notes module 714, and/or the client analysis module 716) may not exist for the user's particular computer platform. This is especially true for the computer platforms used to implement the web clients 304. In these cases, the user will not be able to access any of the enterprise server modules for which the user's platform does not have a corresponding client module. In other embodiments, however, software for all of the client modules are available for all platforms. Accordingly, the user's computer platform is not a consideration with respect to the issue of which enterprise server modules the user is capable of accessing. In these embodiments, the client modules may be implemented using multi-platform enabled software, such as Java or other such software.

In some embodiments, the modules are not loaded on the user's computer. Instead, they are run from a server connected to the user's computer. However, a particular user may not have access to all of the modules on the server, for the reasons discussed above.

Command Dispatch Module

The command dispatch module 406 routes enterprise server API commands received from clients 304, 306 to the enterprise server modules that are responsible for processing the commands. This functionality is represented in step 10110 of FIG. 101, described below. The enterprise server API

-72-

commands are described below. Also described below is the mapping of enterprise server API commands to enterprise server modules (i.e., which enterprise server modules process which enterprise server API commands).

Clients

5 As noted above, the system 302 preferably includes two types of clients, network clients 306 and web clients 304. The network clients 306 and the web clients 304 are discussed in greater detail below.

10 It is noted that the functional capabilities of network clients 306 and web clients 304 may differ in some embodiments, and may be the same in other embodiments (this is described below). However, for simplicity purposes, the present invention is sometimes described as interacting with "clients 304, 306," which may represent one or more network clients 306, one or more web clients 304, or any combination of network clients 306 and web clients 304.

Network Clients

15 The network clients 306 preferably communicate with the enterprise server 314 using the enterprise server 314's natural language, which is called the enterprise server API (application program interface). The network clients 306 can be considered to represent web browsers that are specific to the language and the functions supported by the enterprise server 314. The
20 enterprise server API is described below.

 The network clients 306 preferably communicate directly with the enterprise server 314 via a communication network 312, which is preferably a network that supports the well known HTTP (hypertext transport) protocol.

25 Each network client 306 is preferably implemented using a computer, such as the computer 1102 shown in FIG. 11. Preferably, the computers used

-73-

to implement the network clients 306 are personal computers with at least 16 MBytes of main memory 1108, running the Microsoft Windows 95 operating system or the Microsoft Windows NT operating system. The software executing in the network clients 306 is preferably written in the C++ computer programming language.

In fact, preferably all software of the invention is written using the C++ computer programming language. Database manipulation code is written in a combination of SQL and C++. Other computer programming languages could alternatively be used, such as SmallTalk, Java, C, Pascal, ADA, etc.

Web Clients

The web clients 304 are each preferably implemented using a computer such as that shown in FIG. 11. Commercial web browser software preferably executes in the web clients 304. Any commercial and well known web browser software can be used, such as web browser software from Netscape, Microsoft, Sun, etc.

Unlike the network clients 306, the web clients 304 do not typically utilize the enterprise server API. Accordingly, the web clients 304 communicate indirectly with the enterprise server 314 via the web server 310.

FIG. 8 is a block diagram of the web server 310. This block diagram also illustrates the data flow between the web client 304 and the enterprise server 314 via the web server 310. The web server 310 includes a translator 804 that translates between the respective languages of the web clients 304 (preferably HTML, or hypertext markup language) and the language of the enterprise server 314 (i.e., the enterprise server API).

Specifically, the enterprise server 314 sends raw data 802 to the web server 310 over the network 312. The translator 804 in the web server 310 translates the raw data 802 to data in the well known HTML data format. This

-74-

HTML data 806 is sent to the web client 304 over network 308. A browser 808 in the web client 304 renders the HTML data 806. The translator 804 translates data going from the web client 304 to the enterprise server 314 in a similar manner. It is noted that data formats other than HTML could alternatively be used. In particular, any data format used by the browser 808 could alternatively be used in the invention.

Since the web server 310 communicates with the enterprise server over the network 312 using the enterprise server API, the web server 310 appears to be a network client 306 from the perspective of the enterprise server 314. The interaction between the web clients 304 and the enterprise server 314, and the network clients 306 and the enterprise server 314, is further described below.

The use of commercial web browser software in the web clients 304 is advantageous because such software executes on different computer platforms (that is, there are versions of the browser software that executes on different computer platforms), such as computer platforms produced by IBM, Apple, Sun, SGI, HP, companies producing computers that are compatible with those produced by these companies, etc. Thus, the present invention via the web clients 304 enables users working on any type of computer platform for which a commercial web browser is available to access the enterprise server 314 and the databases 316. This feature of the invention is particularly important in the data processing world of today, where any given corporate entity may have users on many different computer platforms. The present invention allows all of these users (using commercial web browser software, which the corporation may already own) to access and work with the enterprise server 314 and the databases 316, irrespective of the type of computer that they are using.

Enterprise Server API (Application Programming Interface)

The enterprise server API includes commands for accessing functions and capabilities supported by the enterprise server modules (shown in FIG. 4). Many of the enterprise server commands have an implicit parameter which is the sessionID (identifier) of the current server session. A sessionID is obtained by calling the ReqLogin command. All commands issue an exception on failure.

Interaction between the clients 304, 306 and the enterprise server 314 is conducted via direct or indirect use of the enterprise server API, whether or not stated explicitly herein.

Other applications (not discussed herein) may interact with the enterprise server 314 as long as such applications conform to the enterprise server API.

An embodiment of the invention includes timed, automatic executing commands (such as the automatic searches described above). These commands execute upon the occurrence of an event (the event can be defined in the passed parameters of the commands). Such events include time (for example, execute every 30 days), system update (for example, run this search every time new patents are loaded into the enterprise server), data change (for example, automatically regroup these patents every time a corporate entity changes), etc. These timed automatic executing commands are essentially those listed and described below, with additional logic for detecting the occurrence of the defined condition(s)/event(s), and for automatic execution upon such detection.

The commands that make up the enterprise server API according to an embodiment of the invention are described below. It should be understood that the enterprise server API is extendable to support other enterprise server modules, or to support additional or modified functions in existing enterprise server modules, or to support other functions described herein. Embodiments of the enterprise server API may include only a subset of the following

commands. Also, modules other than the ones identified could process the following commands. For example, the Server Administration Module 418 could process the Security Module 402's commands.

*Commands Processed by the Server Administration
Module 418*

5

ReqLogin(username, password)

Returns: sessionID

Description: Login command for the enterprise server 314 for user authentication and to establish a user session on the enterprise server 314.

10

ReqLogout()

Returns: nothing

Description: Terminates user session with enterprise server 314.

ReqAddUsers(userList userSet)

Returns: nothing

15

Description: Adds the users specified in userSet to the system (as being able to access the system).

ReqGetAllUsers()

Returns: list of users in system

Description: Returns a list of users registered to work with the system.

20

ReqGetUsers(userIdList userIDSet)

Description: Returns a list of users (identified by userIDSet) and their user IDs.

ReqRemoveUsers(userIdList userIDSet)

Description: Removes a list of users from system, specified by their user IDs (userIDSet).

Commands Processed by the Security Module 402

ReqGetPermissionList(string docGroupID)

- 5 Description: Gets and returns the permission list for a group specified by docGroupID.

ReqRemovePermission(string docGroupID, string entityID)

Description: Removes all access privileges to a group (specified by docGroupID) from an entity (specified by entityID).

- 10 ReqSetPermission(string docGroupID, string entityID, string mode)

Description: Sets access permission (specified by mode) for an entity (specified by entityID) to use a group (specified by docGroupID).

Commands Processed by the Document Storage and Retrieval Module 408

- 15 ReqCanPage(document, section, page)

Returns: raw image data

Description: Gets the bitmap image (also called the canonical representation) associated with a section and page of a document as specified in the passed parameter.

- 20 ReqDocList()

Description/Returns: Retrieves and returns a list of documents in the repository.

-78-

ReqTxt(document)

Returns: raw text data

Description: Gets the ASCII text of a document specified in the passed parameters.

5 ReqRawCan(document)

Returns: entire image file

Description: Gets the entire collection of images associated with a document specified in the passed parameter.

ReqRawEqv(document)

10 Returns: entire equivalent or text file

Description: Gets the document equivalent data or text data of a document specified in the passed parameter. This return data is a textual representation of the document.

ReqCanHeader(document)

15 Returns: Image file header data

Description: Gets header information about the collection of images associated with the document specified in the passed parameter, including the size, width, and height of the images.

ReqAbstract(spStringSet patentList)

20 Returns: A list of abstract/patent number pairs

Description: Retrieves abstracts associated with the patent list specified in the patentList parameter.

ReqGetAllPatentData(int sindex, int eindex)

Returns: list of patents with their bibliographic information, search handle

-79-

Description: Gets the list of patents plus their bibliographic information, starting from `sindex` and ending with `eindex`, where `sindex` and `eindex` are based on the ordering of the patents in the patent bibliographic databases 604. Example: `ReqGetAllPatentData(0, 5)` returns the first 6 patents in the patent bibliographic databases 604 with their bibliographic information. Also returns a handle that identifies the persistent result set on the enterprise server in order to get more patents from the result set.

`ReqGetAllPatents(int sindex, int eindex)`

Returns: list of patents, search handle

Description: Gets the list of patents, starting from `sindex` and ending with `eindex`, where `sindex` and `eindex` are based on the ordering of the patents in the patent bibliographic databases 604. Example: `ReqGetAllPatents(0, 5)` returns the first 6 patents in the patent bibliographic databases 604. Also returns a handle that identifies the persistent result set on the enterprise server in order to get more patents from the result set.

`ReqGetPatentDataInGroup(string groupID, int sindex, int eindex)`

Returns: list of patents with their bibliographic information, search handle

Description: Gets a list of the patents in a group (specified by `groupID`) with their bibliographic information, starting from `sindex` and ending with `eindex` (relative to the ordering of the patents in the group). Example: `ReqGetPatentDataInGroup(groupid, 0, 5)` returns the first 6 patents and their bibliographic information from the group. Also returns a handle that identifies the persistent result set on the enterprise server in order to get more patents from the result set.

`ReqGetPatents(string groupID, int sindex, int eindex)`

Returns: list of patents from `sindex` to `eindex`.

-80-

Description: This is similar to ReqDocsInGroup (described below), except it returns patents from sindex to eindex from the group specified by groupID. sindex and eindex are relative to the ordering of the patents in the group.

Example: ReqGetPatents(groupid, 0, 5) returns the first 6 patents in the group.

5 ReqGetPatents(groupid, 6, 11) returns the second 6 patents in the group. Also returns a search handle on which to get subsequent patents.

ReqDeletePatentHandle(string handle)

Returns: nothing

10 Description: Deletes a generated result set of patents specified by handle, generated by ReqGetAllPatents or ReqGetPatents.

ReqGetPatentDataFromHandle(string handle, int sindex, int eindex)

Returns: list of patents with their bibliographic information, search handle

15 Description: Receives a handle. From that handle, gets the patents from sindex to eindex (where sindex and eindex are relative to the handle), and also gets their bibliographic information. The handle is generated by ReqGetPatentDataInGroup or ReqGetAllPatentData.

ReqGetPatentsFromHandle(string handle, int sindex, int eindex)

Returns: list of patents, search handle

20 Description: Receives a handle. From that handle, gets the patents from sindex to eindex (where sindex and eindex are relative to the handle). The handle is generated by ReqGetPatents or ReqGetAllPatents

ReqGetPatentsWithBibInfo(patentList list)

Returns: list of patents with their bibliographic information

25 Description: Given a list of patents (specified by list), returns a list of patents with their bibliographic information.

Commands Processed by the Grouping Module 412

- ReqGetGroupHierarchy()
Returns: group hierarchy information
Description: Retrieves information about the hierarchical structure of the groups stored in the databases 316.
- 5
- ReqAddDocGroup(groupParentID, groupID)
Returns: nothing
Description: Adds an existing group (specified by groupID) as a child to another group (specified by groupParentID).
- 10
- ReqRemoveDocGroup(groupParentID, groupID)
Returns: nothing
Description: Unlinks a group (specified by groupID) from its parent group (specified by groupParentID). If the group has no parent, the group is deleted.
- 15
- ReqUpdateDocGroupProperties(group)
Returns: updated group
Description: Update group properties (such as description and title) of the group specified by the passed parameter.
- 20
- ReqNewDocGroup(groupParentID, group)
Returns: new group
Description: Create a new group (corresponding to the group passed parameter) on the enterprise server 314 as a child of another group (specified by groupParentID).

ReqDocsInGroup(groupID)

Returns: list of document names

Description: Get the list of documents in the group specified by the passed parameter.

5 ReqAddDocListToGroup(groupID, documentList)

Returns: nothing

Description: Add a list of existing documents (specified by documentList) to the group specified by groupID.

ReqAddPatents(string groupID, patentList pList)

10 Returns: nothing

Description: Add the list of patents (specified by pList) to the group specified by groupID.

ReqRemoveDocListFromGroup(groupID, documentList)

Returns: nothing

15 Description: Remove a list of documents (specified by documentList) from the group specified by groupID.

ReqRemovePatents(string groupID, patentList pList)

Returns: nothing

20 Description: This is similar to ReqRemoveDocListFromGroup. This removes a list of patents (specified by pList) from the group specified by groupID.

ReqNewGroupWithSearchPatents(parentID, grp, handle)

Returns: group

Description: Creates a new group specified by grp under the parent group specified by parentID with documents/patents from a persistent result set

generated by ReqSearchRelevant. The persistent result set is specified by handle.

Commands Processed by the Notes Module 414

ReqCreateNote(noteID, text, reference)

5 Returns: created note

Description: Create a new note on the enterprise server 314. The identifier of the new note is noteID, the text (or pointer to any type of data in any form) or content of the note is specified by text, and linkage information is specified by reference. The linkage information specifies the document and the portion
10 within the document to which the new note is linked.

ReqUpdateNote(noteID, text, reference)

Returns: updated note

Description: Update a note (specified by noteID) on the enterprise server 314 with new text (specified by text) or reference (i.e., linkage information
15 specified by reference).

ReqAddNoteListToDoc(groupID, documentID, noteList)

Returns: nothing

Description: Add a list of notes (specified to noteList) to a document (specified by documentID) in a group specified by groupID. groupID is used for security
20 purposes (i.e., to ensure that the operator has the proper security level to add the notes to the document).

ReqRemoveNoteListFromDoc(groupID, documentID, noteList)

Returns: nothing

Description: Remove a list of notes (specified by noteList) from a document (specified by documentID) in a group (specified by groupID). groupID is used for security purposes (i.e., to ensure that the operator has the proper security level to remove the notes from the document).

5 ReqNotesOnDoc(groupID, documentID)

Returns: nothing

Description: Get all notes associated with a document (specified by documentID) in a group (specified by groupID). groupID is used for security purposes (i.e., to ensure that the operator has the proper security level to
10 retrieve the notes associated with the document).

AddGroupNote(groupID, gnote)

Returns: A group note

Description: Adds a new group note represented by gnote to the group identified by groupID. Updates the note if it already exists.

15 AddPatentNote(groupID, note)

Returns: note

Description: Adds a new patent note represented by the parameter note to the group identified by groupID. Updates the note if it already exists.

GetGroupNotes(groupID)

20 Returns: returns group notes

Description: Retrieves all group notes associated with the group identified by groupID.

GetPatentNotes(groupID)

Returns: returns patent notes

Description: Retrieves all patent notes associated with the group identified by groupID.

RemoveGroupNote(groupID, groupNoteID)

Returns: nothing

5 Description: Removes the group note with groupNoteID from the group specified by groupID.

RemovePatentNote(groupID, noteID)

Returns: nothing

10 Description: Removes the patent note with NoteID from the group specified by groupID.

UpdateGroupNote(gnote)

Returns: a group note

Description: Updates an existing group note.

UpdatePatentNote(note)

15 Returns: a patent note

Description: Updates the properties of an existing note.

AddNoteSegment(noteID, noteseg)

Returns: a note segment

20 Description: Add given note segment represented by noteseg to a patent note identified by noteID.

GetGroupNotesMatchingString(groupID, search)

Returns: List of notes

-86-

Description: Return group note identifiers of notes in a group (specified by groupID) that contain the search string (represented by the search parameter).

GetNoteSegments(noteID)

5 Return/Description: Get and return all note segments associated with the given patent note specified by noteID. Also return their location information.

GetNoteSegmentsMatchingString(groupID, search)

Returns: List of note segments

Description: Return note segment identifiers of notes in a group (specified by groupID) that contain the search string (represented by the search parameter).

10 GetPatentLocations(groupID, patentName)

Returns: patent location list

Description: Returns all patent locations attached to the patentName.

LinkNoteSegment(noteSegmentID, location)

Returns: nothing

15 Description: Links note segment to location in a patent.

UnlinkNoteSegment(noteSegmentID)

Returns: nothing

Description: Unlinks note segment from location in patent.

UpdateNoteSegment(noteseg)

20 Returns: note segment

Description: Updates an existing note segment.

RemoveNoteSegment(noteID, noteSegmentID)

Returns: nothing

Description: Removes note segment specified by noteSegmentID from a patent note specified by noteID.

Commands Processed by the Searching Module 410

5 ReqSearch(searchParameters, startIndex, endIndex)

Returns: list of search results, search handle

Description: Execute a search based on searchParameters, retrieve search results from startIndex to endIndex in result table.

10 ReqRetrieveSearchResult(searchHandle, startIndex, endIndex)

Returns: list of documents

Description: Retrieve search results of previously executed search (identified by searchHandle) from startIndex to endIndex in result table. Also implemented as ReqRetrieveSearchRelevantResult(string handle, int sindex, int eindex).

15 ReqSearchHighlights(searchHandle, documentID)

Returns: list of text offsets for highlighting

Description: Retrieve search hit information for a particular document (specified by documentID) so that search term highlighting can be performed on the document. The search is specified by searchHandle.

20 ReqSearchBib(spSearchParameters s, int sindex, int eindex)

Returns: list of search results, search handle

Description: Executes a search based on SearchParameters, retrieves search results from startIndex to endIndex in result table, where the search results include the bibliographic information of the documents identified by the search.

-88-

ReqSearchRelevant(ReqRetrieveSearchRelevantResult(searchType,
searchOrder, query, sindex, eindex, minRelevance)

Return: search results with bibliographic information

Description: Performs a search on either the repository patents, all patents, or
5 patents not in the repository (selected by searchType) using the search
parameters. Returns the results sorted by field specified in searchOrder. Only
return results that have a relevance number greater than minRelevance. Gets
the results from row sindex to row eindex. Also returns a handle that identifies
the persistent result set on the enterprise server in order to get more patents from
10 the result set.

ReqRetrieveSearchRelevantResult(handle, sindex, eindex)

Returns: List of patent bibliographic information, and a search handle

Description: Retrieves portions of a persistent search result set generated by
ReqSearchRelevant and identified by handle, from row sindex to row eindex.

15

Commands Processed by the Analysis Modules 416

ReqFunction(function, level, PatExp1, PatExp2, PatTerm1, PatTerm2,
GroupID)

Returns: Results generated by performing the function specified by the function
passed parameter

20

Description: Performs the function specified by the function passed parameter.
The function passed parameter can identify any function performed by the
enterprise server 314, such as Patent Mapping, Patent Aging, Patent Citation,
Inventor Employment Information, and Patent Count. Level specifies the
number of levels in the group hierarchy to drill down. PatExp1 and PatExp2
25 specify two dates (a time range) that designate a search scope based on patent
expiration (see FIG. 128). PatTerm1 and PatTerm2 specify two dates (a time

range) that designate a search scope based on patent term remaining (see FIG. 128). GroupID identifies a group for which the function is performed.

ReqAnalysisCitation (documentID, direction, levels)

Returns: patent/child table and patent bibliographic information table

5 Description: Performs a patent citation function for a patent identified by documentID. Either a forward or a backward citation function is performed, as identified by the direction parameter. A multilevel citation function can be performed. The number of levels is specified by the levels parameter. For a detailed description of the operation of the enterprise server 314 when
10 processing the ReqAnalysisCitation command, see step 15806 (FIG. 158), discussed below.

Client/Server Interaction

FIG. 101 is a flowchart depicting the generic interaction between the network clients 306 and the web clients 304 with the enterprise server 314. In
15 step 10106, a client 304, 306 (either a network client 306 or a web client 304) sends a command and possibly information to the Enterprise server 314. The command represents a request to a module of the Enterprise server 314 to perform some function (which is called the requested function for reference purposes). In the case of the web client 304, the web server 310 translates the command to a
20 message that conforms with the enterprise server API. In the case of the network client 306, the command already conforms to the enterprise server API.

In step 10108, the security module 402 and the server configuration module 404 (see FIG. 4) in the Enterprise Server 314 determine whether or not the client 304, 306, or the user who is using the client 304, 306, is authorized to
25 issue the command. More particularly, the security module 402 determines whether or not the client 304, 306 or user has the appropriate security privileges

-90-

with respect to the target information in the databases 316 (that is, the information in the databases 316 that will be accessed if the requested function is performed by the enterprise server 314). The server configuration module 404 determines whether or not the client 304, 306 or user has access to the requested function (i.e., has the client 304, 306 or user been configured to access the requested function?). If the client 304, 306 or user is authorized to issue the command, then step 10110 is performed. Otherwise, the enterprise server 10110 does not honor (does not perform) the command (an appropriate message or exception may be sent back to the client 304, 306).

In step 10110, the command dispatch module 406 in the enterprise server 314 routes the command and the information to the appropriate enterprise server module. These enterprise server modules are described above.

In step 10112, the enterprise server module performs the requested function.

In step 10114, the Enterprise server 314 sends the results of the command to the client 304, 306. In the case of the web client 304, the web server 310 translates the results received from the enterprise server 314 to the language supported by the web client 304 (preferably HTML), and then forwards these translated results to the web client 304 over the network 308. This network may be a local area network or a wide area network (it can even be a global network).

In step 10116, modules in the client 304, 306 operate to display the results to the user, and also operate to enable the user to manipulate, process, and otherwise utilize the results.

FIG. 7 illustrates client modules 701 of a web client 304 and a network client 306 according to an embodiment of the invention. The modules in a web client 304 or a network client 306 correspond to modules in the enterprise server 314. A given module in the enterprise server 314 performs specific tasks pursuant to commands received from the clients 304, 306. The performance of these tasks generates results. For example, the document storage and retrieval module 408

in the enterprise server 314 stores and retrieves information from the databases 316.

The corresponding module in the clients 304, 306 receives these results and presents them to the user, and enables the user to work with and manipulate the results. For example, the client document storage and retrieval module 708 presents to the user the information retrieved by the document storage and retrieval module 408 in the enterprise server 314.

The process just described is an iterative one, as represented by control arrow 10118.

The interaction between the web clients 304 and the enterprise server 314, and the network clients 306 and the enterprise server 314, is further described with reference to FIGS. 81-83. FIG. 81 generically depicts the interaction between the enterprise server 314 and the clients 304, 306. A client 304, 306 (either a web client 304 or a network client 306) sends, for example, a ReqSearch command to the enterprise server 314. The searching module 410 in the enterprise server 314 processes the ReqSearch command, and returns a handle (pointer) to the search results plus a list of documents from startIndex to endIndex. The client 304, 306 then issues a ReqRetrieveSearchResult in order to obtain the search results. The searching module 410 processes the ReqRetrieveSearchResult, and returns a list of documents.

FIGS. 82 and 83 contrasts, in greater detail, the interaction between the enterprise server 314 and the network clients 306, and the enterprise server 314 and the web clients 304. FIG. 82 represents the interaction between the enterprise server 314 and a network client 306. The network client 306 sends a ReqRawCan command to the enterprise server 314. The document storage and retrieval module 408 processes this command, and returns image data representative of the document. Since the network client 306 supports the enterprise server API, translation of the command or the return result is not necessary.

FIG. 83 represents the interaction between the enterprise server 314 and a web clients 304. The web client 304 issues a command to retrieve a document image. This command may or may not conform to the enterprise server API. If it does not conform, the web server 310 translates the command to one that conforms to the enterprise server API. The enterprise server 314 returns an image. However, the image is not in a format that the web client 304 supports. Thus, the web server 310 converts the image to a format that the web client 304 supports. In the example of FIG. 83, the web server 310 converts the image to the well known JPEG format, then forwards the converted image to the web client 304. Other image formats could also be used, and depends on the formats used by the commercial web browser software being used in the web clients 304.

Patent-Centric URL Commands

The interaction between the enterprise server 314 and the web clients 304 shall now be more particularly described.

FIG. 151 illustrates the interaction between the Enterprise server 314 and the web client 304. FIG. 151 is similar to FIG. 8. However, FIG. 151 focuses on the commands that are transferred between the web client 304 and the Enterprise server 314, whereas FIG. 8 focuses on the data which is transferred between the web client 304 and the Enterprise server 314.

As discussed above, the web client 304 preferably includes a browser 808, which can be any commercially available browser, such as (but not limited to) those available from Netscape, Microsoft, IBM, SUN, Novell, etc. In an embodiment of the invention, the browser 808 issues URL (Uniform Resource Locator) commands in order to access and retrieve data from the databases 316 via the enterprise server 314.

The general format of URL commands is well known, and is presented in FIG. 150 for the convenience of the reader. A URL command 15002 includes a protocol field 15004, a destination field 15006, and a command field 15008.

5 The protocol field 15004 specifies the protocol that is to be used in transporting the URL command 15002 from its source to its destination. Example protocols include HTTP and FTP (file transfer protocol).

10 The destination field 15006 specifies the destination of the URL command 15002. For example, the destination field 15006 may include information that identifies the server (such as the enterprise server 314) from whom information is being requested from.

The Command field 15008 stores information representing a command or an action or an identification of requested data. The effect of the URL command 15002 is to request that the entity identified by the destination field 15006 perform the command or action specified in the command field 15008.

15 According to the invention, the commands inserted into the command field 15008 are patent-centric or patent-specific. Accordingly, the URL commands generated by the present invention are patent-centric, or patent-specific.

20 In practice, HTML data representative of a web page is transferred to the browser 808 in the web client 304 upon connection with the Enterprise server 314. The browser 808 processes this received HTML data and, as a result, displays a web page. HTML processing is well known. Examples of web pages are shown in FIGS. 53 and 140. An operator at the web client 304 enters information into the fields of the web page.

25 For example, with reference to FIG. 53, an operator could enter an inventor name into the inventor field 5310 of web page 5302. The information entered by the operator represents a search string. After receiving an appropriate user command (which the user issues by pressing the Search button 5334, for example), the browser 808 in accordance with the software associated with the displayed web page 5302 generates one or more URL commands. These URL

commands include the information entered by the operator into the fields of the displayed web page 5302. In other words, these URL commands include the search string entered by the operator. These URL commands are directed to the enterprise server 314, and request the enterprise server 314 to conduct one or more searches of the databases 316 in accordance with the operator-supplied search string.

Referring again to FIG. 151, a translator 804 in the Web server 310 receives the URL commands from the browser 808 in the web client 304. The translator 804 converts the URL commands to Enterprise server API commands. These Enterprise server API commands are received and processed by the Enterprise server 314. The Enterprise server API commands are discussed above.

According to an embodiment of the present invention, the URL commands sent from the browser 808 in the web client 304 to the translator 804 in the Web server 310 conform to a patent-centric URL command language. Accordingly, the URL commands sent from the browser 808 to the translator 804 represent patent-centric URL commands.

The patent-centric URL command language of the present invention essentially represents an API (Application Programming Interface) of the Web server 310. The patent-centric URL command language of the present invention includes the following patent-specific commands that are inserted in the command field 15008 of URL commands. It should be understood that the following is a representation of the types of commands that can be placed into the command field 15008 of URL commands according to the invention. The invention can support other patent-centric/specific commands. Thus, the following is provided for purposes of illustration, not limitation.

-95-

Command: Search

Description: This command is used to instruct the enterprise server 314 to perform a search or to retrieve more results from a previous search.

5	Parameters:	scope	Identifies the patents or documents that should be searched, such as all patent/documents in bibliographic databases 602, all patents/documents in the document databases, etc.
10		number	Specific patent (document) number to search for, if the operator specified such a search
		title	Specific patent/document title to search for, if the operator specified such a search
		inventor	Specific inventor name to search for, if the operator specified such a search
15		assignee	Specific assignee name to search for, if the operator specified such a search
		class	USPTO class/subclass to search for, if the operator specified such a search
20		udk	user defined keyword to search for, if the operator specified such a search
		datesearchtype	Equals 1 if date search criteria is BEFORE, AFTER, or ON; Equals 2 if date search criteria is BETWEEN (if the operator specified a date search)
25		datequalifier	set equal to one of : AFTER, BEFORE, or ON (i.e., modifies the datesearchtype when datesearchtype is equal to 1) (if the operator specified a date search)

-96-

	date1	date for BEFORE, AFTER, or ON (if the operator specified a date search)
	date2	first date of BETWEEN (if the operator specified a date search)
5	date3	second date of BETWEEN (if the operator specified a date search)
	abstractquery	search string for abstract, if the operator specified such a search
10	fulltextquery	search string for full text, if the operator specified such a search
	showtitle	flag to show title in search results
	showdate	flag to show date in search results
	showinventor	flag to show inventor in search results
	showassignee	flag to show assignee in search results
15	showudk	flag to show udk (user defined keyword) in search results
	orderby	Specifies ordering preference; one of: RELEVANCE, PATENT NUMBER, etc.
20	guid	contains the GUID (globally unique identifier) of a search results table (generated by a prior search) that is potentially still on the enterprise server 314
	begin	beginning part of a hitlist range displayed to the user
	end	ending part of the hitlist range displayed to the user
25	rel	a generic "catchall" field. Used to deal with characteristics of HTML. Possible values include: NEXT, PREVIOUS, FIRST, LAST, GET_RESULTS_IN_FILE, and SKIM_IMAGES.
	total	total number of hits

numberpage number of hits to be displayed per page

Operation:

An operator defines a search using, for example, the patent search screen 14002 of FIG. 140. The web client 304 generates a patent-centric URL command that contains a search command in the command field 15008. The parameters of the search command reflect the search defined by the operator. For example, if the operator defined a search based on patent number and PTO class, then the number and class parameters of the search command would be filled in.

For a new search, the begin parameter is equal to 0, indicating that the enterprise server 314 should return the search results beginning from record 0 (of the search results). The numberpage parameter indicates to the enterprise server 314 the number of items to return search results on. For example, if the numberpage parameter equals 10, then the enterprise server 314 returns information on 10 search results items, starting with the first item (where begin is equal to 0).

In a new search, the total parameter is returned by the enterprise server 314, and represents the number of search hits. This value of the total parameter is then returned to the enterprise server 314 in any subsequent, related search commands (to obtain additional search results, for example). The browser 808 receives the search results provided by the enterprise server 314 and displays the search results in a screen such as that shown in FIG. 141. The operator presses the navigation icons 14108 to obtain additional search results. If the operator presses a Next icon 14109, for example, a new URL command with the search command in the command field 15008 is generated to obtain the next page of search results.

For this and any other subsequent, related search commands, the parameters (such as scope, search key terms, items to display, etc.) that define the search are the same as in the original search command. In this and subsequent related search commands, however, the GUID of the original search results is

provided. The enterprise server 314 uses this GUID to access the original search results, if they are still available on the enterprise server 314 (they may have expired). If the original search results are not still available on the enterprise server 314, then the enterprise server 314 reexecutes the search.

5 In the case where the operator pressed the Next icon 14109, the rel field is set equal to NEXT. If the operator pressed the Previous icon 14111, the rel field is set equal to PREVIOUS. If the operator presses a Last icon 14119, the rel field is set equal to LAST. If the operator pressed a First icon 14191, the rel field is set equal to FIRST.

10 In the case where the operator pressed the Next icon 14109, the begin field is set to 0. The enterprise server 314, upon receipt of this command, identifies the next page of search results to send to the web client 304. The enterprise server 314 does this by adding the value of the numberpage to the value of begin. The result of this addition operation identifies the first record in the search results to
15 send to the web client 304. Starting from this value, the enterprise server 314 sends numberpage items to the web client 304. Thus, where the begin field is equal to 0, and the numberpage is equal to 10, the enterprise server 314 sends records 10-19 of the search results to the web client 304.

20 Suppose that the operator again presses the Next icon 14109. In the resulting search command, the begin field is set to 10 (this is done by software associated with the link associated with the Next icon 14109 in the display 14102 in FIG. 141). The enterprise server 314, upon receipt of this command, sends numberpage items to the web client 304 starting from begin plus numberpage. Thus, where the begin field is equal to 10, and the numberpage is equal to 10, the
25 enterprise server 314 sends records 20-29 of the search results to the web client 304.

 If rel is equal to GET_RESULTS_IN_FILE, then the enterprise server 314 returns HTML data to the web client 304. The web client 304 prompts the

-99-

operator for a file name, and then saves this data in the file identified by the file name.

If rel is equal to SKIM_IMAGES, then the enterprise server 314 returns information that identifies two frames (or windows or panes). The enterprise server 314 supplies, for one of the frames, information that identifies a file in the enterprise server 314. This file stores a list of the search results. The web client 304 issues a URL command to retrieve this file from the enterprise server 314. The web client 304 then displays the list of the search results into this first frame. For the second frame, the enterprise server 314 provides information that identifies the location of the image of the first item in the search results. The web client 304 issues a URL command to retrieve this image (or a portion of the image, like the portion of the image corresponding to the first page of the document) from the enterprise server 314. The web client 304 then displays this image in the second frame. This two frame display is shown, for example, in FIG. 148.

In the above description, the enterprise server 314 is described as interpreting the URL commands generated by the web client 304. In practice, according to a preferred embodiment of the invention, the translator 804 in the web server 310 translates the URL commands to commands in the enterprise server API language. The enterprise server 314 then processes these enterprise server API commands as discussed herein. Such translation is described below.

Command: GetText

Description: This command instructs the enterprise server 314 to retrieve and return the text of a patent. Preferably, the web server 310 returns the entire text of the patent to the web client 304. However, the following additional parameters are used for text/image synchronization purposes (i.e., when the operator switches between displaying text and displaying images).

-100-

Parameters: section Image files can contain multiple sections. The first section is the image of a patent. Subsequent sections contain post-issuance documents, such as a certificate of correction. This parameter describes the specific section desired from the image file and is used to help synchronize text and image display. Specifically, suppose that the operator is viewing an image. Then the operator elects to view text, such that a GetText command is generated. In the section and page parameters of this GetText command, the section and page of the image being viewed (when the GetText command is generated) is stored. Thus, if the operator then elects to switch back to viewing the image, the invention can identify which image to display by reference to the values stored in these section and page parameters.

page Sections are decomposed into pages. This parameter keeps track of the page.

currentview One of SPLIT_SCREEN or FULL_SCREEN. This parameter identifies the current display mode, and is used by the invention to identify icons that should be active and inactive. For example, if currentview is equal to SPLIT_SCREEN, then any split screen icons are greyed out, to indicate that they are inactive.

number Identifies the document to retrieve (such as a patent number)

-101-

Command: GetImage

Description: This command instructs the enterprise server 314 to retrieve and return the image of a patent. Parameters are the same as described for GetText: section, page, and currentview. Section and page identify which image to retrieve.

5

Command: GetSplitScreen

Description: This command returns two views of a patent, one of the text and one of the image. Number identifies the document whose text is retrieved and displayed. Section and page identify the image which is retrieved and displayed. In an embodiment, the enterprise server 314 returns information on two side-by-side frames (or panes or windows). For one of the windows, the enterprise server 314 provides a GetText command (or equivalent enterprise server API command, which is then translated to the GetText command), to return the text corresponding to the document identified by number. In the other window, the enterprise server 314 provides a GetImage command (or equivalent enterprise server API command, which is then translated to the GetImage command), to return the image corresponding to section and page. The web client 304 then executes these GetText and GetImage command to retrieve this information from the enterprise server 314.

10

15

20

Parameters	number	The number of the patent
	section	same as GetText
	page	same as GetText
	currentview	same as GetText

25

Command: GetAbstract

-102-

Description: This command instructs the enterprise server 314 to return the abstract and predefined bibliographic fields of the patent (or other document) specified by the number parameter.

Parameters: number Patent (document) number

5 Command: GetTextOrAbstract

Description: This command instructs the enterprise server 314 to return the text of a patent as specified by the number parameter, if available (if in the document databases 612); otherwise, it returns the abstract of the patent and some other predefined bibliographic fields.

10 Parameters: number Patent (document) number

Command: OrderPatents

Description: This command, when executed, generates a message representing an order form to order electronic copies of the patents specified in the parameter list. This message is then sent to a third party provider, or sent to a party within the customer corporate entity who will take the message and then order electronic copies of the patents.

15

Parameters: item0 first patent in the list
item1 second patent in the list
20 item2 third patent in the list
...
itemN Last item in the list.

20

Command: GetCitationTree

Description: This command instructs the enterprise server 314 to perform a patent citation function for a patent identified by the number parameter. Either a backward or forward function is performed,

25

-103-

as indicated by the direction parameter. The number of levels for the citation function is specified by the levels parameter. GetCitationTree maps to ReqAnalysisCitation, described above. For a detailed description of the operation of the enterprise server 314 when processing the ReqAnalysisCitation command, see step 15806 (FIG. 158), discussed below.

Parameters: number Identifies a patent
 direction Identifies either a backward or forward citation
 function
 levels Identifies the number of levels for the citation
 function.

Translation

As described above, the translator 804 in the web server 310 translates between patent-centric URL commands and enterprise server API commands (see FIG. 151). Such translation according to an embodiment of the invention is shown in the following table. Other translations between patent-centric URL commands and enterprise server API commands are possible.

Table 4

Patent-centric URL command language	Enterprise Server API language
search	ReqSearchRelevant or ReqSearch
GetText	ReqTxt
GetImage	ReqCanPage

-104-

Patent-centric URL command language	Enterprise Server API language
GetSplitScreen	Returns data that includes representations of a GetText command and a GetImage command (as described above). GetText corresponds to ReqTxt, and GetImage corresponds to ReqCanPage.
GetAbstract	ReqSearchBib
GetTextOrAbstract	In the web client or the translator, this command is translated to GetText (corresponding to ReqTxt). If GetText is not successful, then GetAbstract is issued (corresponding to ReqSearchBib). See FIG. 152.
OrderPatents	None
GetCitationTree	ReqAnalysisCitation

5

Client Architecture

10

FIG. 114 illustrates a logical representation of the architecture of the network client 306 according to an embodiment of the invention. This architecture also applies to the web client 304 in some embodiments. In other embodiments, the web client 304 is represented by well known browser software executing in a computer, such as that shown in FIG. 10.

-105-

The User Interface layer 11404 may be based on various well known user interfaces. The user interface layer in the web client 304 preferably processes HTML data. The user interface layer 11404 in the network client 306 preferably uses MFC, or the Microsoft Foundation Class Library. In some embodiments, the user interface layer 11404 is built using multi-platform enabled languages, such as Java.

The client modules 701 are shown in FIG. 7, and are described elsewhere herein.

The transaction management layer 11406 implements the specific business-related operations of the invention. These operations include creating a group or changing the security permissions of the group (this could alternatively be done by the domain layer 11408). The transaction management layer 11406 interacts with the client modules 701 to perform these functions.

The domain layer 11408 includes all of the objects that are required to properly implement a patent-centric decision support system. These objects in the domain layer 11408 represent the business and other high level intelligence of the client 304, 306, and enables the client 304, 306 to work with business rules, notes, analysis modules, etc.

Supporting the domain layer 11408 is a broker layer 11410 that provides for sophisticated brokering and caching of objects in the client 304, 306. The broker layer 11410 is responsible for managing the communication between the domain layer 11408 and the server layer 11416. This decouples the domain layer 11408 from the enterprise server 314 and provides for maximum flexibility in the implementation of different enterprise servers 314.

The caching subsystem 11412 of the broker layer 11410 provides a means for objects to be cached on the client 304, 306 after they have been retrieved from the enterprise server 314. The caching subsystem 11412 enables the client 304, 306 to manage an infinite number of objects obtained from the enterprise server 314 by only storing those objects that have been most recently used. In an

-106-

embodiment of the invention, the client 304, 306 utilizes a demand paging algorithm. In an embodiment of the invention, caching only takes place on the network client 306.

5 A demand paging algorithm according to an embodiment of the invention is represented by a flowchart 14902 shown in FIG. 149. Preferably, the present invention utilizes a two-level demand paging methodology. The first level of the demand paging methodology is performed by the Caching subsystem 11412 in the client 306, 304, and is represented by steps 14950 in flowchart 14902. The second level of the demand paging methodology is performed by the Enterprise
10 server 314, and is represented by steps 14952 in the flowchart 14902.

In step 14906, the Cache subsystem 11412 receives a request for data from a requester (the requester is typically an upper layer in the architecture). This data request is described herein as being a request for patent data. However, the discussion described herein applies to both patent and non-patent data.

15 At this point, it would be useful to describe an example of how a data request is generated. Consider an example console user interface 11802 shown in FIG. 118 (FIG. 118 is further described below). A group hierarchy 11712 is shown in a first pane 11704 of the console user interface 11802. Patents contained within a selected group of the group hierarchy 11712 are listed (with their bibliographical information) in a second pane 11706. In the example shown
20 in FIG. 118, a repository group 11710 is selected. Accordingly, the patents in the repository group 11710 are listed (along with their bibliographical information) in the second pane 11706 of the console user interface 11802.

In a preferred embodiment of the invention, at any given time, all of the
25 information pertaining to the patents in the repository 11710 is not stored in the client 304, 306. Instead, only a portion of the information pertaining to the patents in the repository 11710 are stored in the client 304,306. Preferably, the client 304, 306 retrieves data from the databases 316 as it needs it.

For example, 13 patents are currently listed in the second pane 11706. The client 304, 306 need only store information on these 13 listed patents. In practice, however, at any given time the client 304, 306 may store information on more than the patents being displayed in the second pane 11706. More generally, at any given time, the client 304, 306 may store information on more than the patents being processed, analyzed, displayed, etc., at the client 304, 306. However, the client 304, 306 does not store all of the data from the repository 11710 (unless, for some reason, the client 304, 306 is processing all of that data). Also, during a session with the enterprise server 314, the client 304, 306 does not discard information that it has received from the enterprise server 314, even when the information is no longer being used at the client 304, 306. In other embodiments, the client 304, 306 discards unused data received from the enterprise server 314 in order to make room for additional data.

The client 304,306 retrieves information from the enterprise server 314 by sending a data retrieval request to the Caching subsystem 11412. The Caching subsystem 11412 receives this request in step 14906.

Further in step 14906, the Caching subsystem 11412 identifies from the data retrieval request the patent and the portions of the patent that are being requested. Preferably, in the present invention, a patent has multiple parts or portions. These parts include the patent bibliographic information, the patent (equivalent) text file, and the patent image file. A given data retrieval request may be requesting any or all of these portions of the patent. In step 14906, the Caching subsystem 11412 identifies from the data request which patent is being requested, and also which portions of the patent are being requested. For reference purposes, the patent that is being requested is called the identified patent, and the portions of the identified patent that are being requested are called the identified portions.

In step 14910, the Caching subsystem 11412 determines whether the identified portions of the identified patent are already stored in the local cache.

-108-

For purposes of the present invention, the local cache is represented by the main memory 1108 in the client 304, 306. In alternate embodiments of the invention, the local cache is represented by cache memory in the client 304,306.

5 If the identified portions of the identified patent are currently in the local cache, then in step 14912 the Caching subsystem 11412 retrieves those identified portions of the identified patent from the local cache and returns them to the requester. Operation of flowchart 14902 is then complete, as indicated by step 14922.

10 If, in step 14910, the Caching subsystem 11412 determines that the identified portions of the identified patent are not in the local cache, then step 14914 is performed. In step 14914, the caching subsystem 11412 sends a message to the Enterprise server 314 to request retrieval of the identified portions of the identified patent from the databases 316.

15 As should be clear by the above description of steps 14906, 14910, 14912, and 14914, the caching subsystem 11412 in the client 304,306 operates according to a caching methodology in which some data is stored in the local cache. When retrieving data, the Caching subsystem 11412 first looks in the local cache to determine whether the requested data is located in the local cache. If the data is not found in the local cache, then the Caching subsystem 11412 requests the data
20 from the Enterprise server 314.

This caching methodology performed by the Caching subsystem 11412 in the client 304,306 represents a first level caching methodology according to the present invention. As mentioned above, the Enterprise server 314 performs a second level caching methodology. This second level caching methodology shall
25 now be described.

In step 14916, the Enterprise server 314 receives the message from the client 304,306. The Enterprise server 314 determines whether the identified portions of the identified patent (as indicated in the received message) are currently stored in the local cache of the Enterprise server 314. The local cache

-109-

of the Enterprise server 314 is represented by the Main memory 1108 of the Enterprise server 314. In an alternative embodiment, the local cache of the Enterprise server 314 is represented by Cache memory in the Enterprise server 314.

5 If the identified portions of the identified patent are located in the local cache of the Enterprise server 314, then in step 14920 the Enterprise server 314 retrieves the identified portions of the identified patent from its local cache, and sends this retrieved data to the client 304,306.

10 If, in step 14916, the Enterprise server 314 determines that the identified portions of the identified patent were not located in its local cache, then in step 14918 the Enterprise server 314 retrieves the identified portions of the identified patent from the databases 316. The Enterprise server 314 then in step 14920 returns this retrieved data to the client 304,306. The operation of flowchart 14902 is complete after the performance of step 14920, as indicated by step 14922.

15 As described above, in an embodiment of the invention, the Enterprise server 314 retrieves and returns only the identified portions of the identified patent. In some cases, the Enterprise server 314 instead returns data representative of a plurality of patents, where such data includes the identified portions of the identified patent. This is called the bulk or cluster retrieval mode of the invention.

20 Consider again FIG. 118. As described above, as the operator scrolls through the patents listed in the second panel 11706 of the console 11802, the caching subsystem 11412 sends requests to the Enterprise server 314 to retrieve additional patent data for display in the second panel 11706. When responding to such requests involving the console 11802, the Enterprise server 314 preferably returns patent data representative of a plurality of patents. Specifically, the Enterprise server 314 returns data representative of a patent cluster.

-110-

A patent cluster represents a given number of patents. In an embodiment of the invention, the number of patents in a patent cluster is equal to 50, but this value is tuneable, and this value may be different for different contexts of the invention. When operating according to the cluster or bulk mode, the Enterprise server 314 in step 14916 determines whether the identified portions of the identified patent are in a cluster that is stored in the local cache of the Enterprise server 314. If the identified portions of the identified patent are in a cluster stored in the Enterprise server 314, then in step 14920 the Enterprise server 314 sends data representative of this cluster of patents to the client 304, 306. If, instead, the Enterprise server 314 determines in step 14916 that the identified portions of the identified patent are not in a cluster currently stored in the Enterprise server 314, then the Enterprise server in step 14918 retrieves data representative of the cluster from the databases 316. The Enterprise server 314 then, in step 14920, sends this retrieved information to the client 304, 306.

Supporting both the domain layer 11408 and the broker layer 11410 is the persistence mapping layer 11414. The persistence layer 11414 is responsible for managing all interactions between the domain layer 11408 and a specific persistent storage device (not shown in FIG. 114). This decouples the domain objects from a particular physical representation, and enables various performance optimizations to be made.

Supporting all aspects of the client 304, 306 is the abstract server interface 11416. This layer 11416 presents a logical server to the other layers of the client 304, 306. It is important to note that all interactions between the client 304, 306 and the enterprise server 314 take place using a high-level, patent-centric set of business decision system command objects. These command objects represent atomic transactions between the client 304, 306 and the enterprise server 314. These commands (or command objects) allow the client 304, 306 to communicate with the enterprise server 314 in a manner that decouples the client 304, 306 from any specific physical implementation of the enterprise server 314. For example, the

-111-

enterprise server 314 could be running MS-DOS and storing objects in a flat file, or be running Unix and storing objects in Informix. As long as the enterprise server 314 responds to the set of command requests presented by the client 304, 306 through the abstract server interface 11416, the client 304, 306 will work correctly. These commands represent the Enterprise server API commands.

The abstract server interface 11416 communicates to make a physical connection to the enterprise server 314. This is done through the network transport layer 11418, which is responsible for taking command objects and transmitting these command objects over a suitable communications network. The network transport layer 11418 also manages appropriate context information that is needed to manage a network connection. The network transport layer 11418: (a) does not require the physical presence of a network - it is possible to run the client 304, 306 and enterprise server 314 on the same physical machine and still have the system used properly; and (b) does not require the use of any specific network, even though one implementation of the system will be based on HTTP (HyperText Transport Protocol).

Additionally, general features of the architecture of FIG. 114 are described in Luke Hohmann, *Journey of the Software Professional A Sociology of Software Development*, Prentice Hall PTR, New Jersey, 1997, which is incorporated herein by reference in its entirety.

Databases

Referring to FIG. 6, some of the databases 316 are described in detail below. In particular, the document bibliographic databases 602, the group databases 621, the person databases 632, the employee databases 634, the security databases 636, the financial databases 638, and the methodology support databases 642 are described in detail below. Both the database structure and the methodology for extract and load of these databases are described below. The

-112-

document databases 612, in particular the patent database 614, and the notes
databases 640 are not described below since they are thoroughly covered in U.S.
Patent No. 5,623,681, U.S. Patent No. 5,623,679, pending application Ser. No.
08/341,129, and pending application Ser. No. 08/590,082, all of which are
5 incorporated by reference herein.

The database structures of the document bibliographic databases 602, the
group databases 621, the person databases 632, the employee databases 634, the
security databases 636, the financial databases 638, and the methodology support
databases 642 are shown in FIGS. 12B-12M. These figures also depict the
10 interaction and connection between these databases. FIG. 12A illustrates the
preferred orientation of FIGS. 12B-12M with respect to one another.

It should be understood that the tables and attributes shown in FIGS. 12B-
12M only represent one embodiment of the present invention. The data in the
databases 316 could be stored using other combinations of tables and attributes.
15 Such other combinations of tables and attributes will be apparent to persons
skilled in the relevant arts based on the discussion contained herein. Accordingly,
the tables and attributes are shown in FIGS. 12B-12M only for purposes of
illustration, and not limitation.

FIG. 45 is a generic dataflow diagram illustrating the general manner in
20 which the databases 316 are loaded with data in accordance with an embodiment
of the invention. FIG. 95 is a flowchart 9502 representing the operation of this
general extract and load procedure. In practice, the initial extract and load of the
databases 316, and/or the updating of the databases 316, may be performed by
employees of the customer and/or consultants retained by the customer.

25 In step 9506, the customer provides data 4504 for upload into the
databases 4508 being processed. The customer provided data 4504 is pertinent
to the databases 4508. For example, if the databases 4508 are intended to store
financial information, then the customer provided data 4504 would comprise

-113-

financial information of interest to the customer, including possibly both the customer's financial information and financial information of competitors.

In step 9508, a filter 4506 modifies the format of the customer provided data 4504 to conform to the database format of the databases 4508. The structure and operation of database filters, such as filter 4506, are well known.

In step 9510, the formatted customer provided data is loaded into the databases 4508. More particularly, the formatted customer provided data 4504 is loaded into a portion of the databases 4508, called a first part 4510 of the databases 4508. Remaining portions of the databases 4508, called the second part 4512 of the databases 4508, cannot be loaded using only the formatted customer provided data 4504. Instead, loading of the second part 4512 may require other information 4516 pertinent to the databases 4508. Additionally, loading of the second part 4512 may require analysis of such additional information 4516 in conjunction with the information in the first part 4510 of the databases 4508. Such analysis is performed by operators 4514 with, potentially, the assistance of the system of the invention.

Accordingly, in step 9512, other information 4516 pertinent to the databases 4508 is obtained.

In step 9514, methodology reports 4518 are run, as needed. Such methodology reports 4518 represent the result of automatic processing and analysis of the databases 4508 with other tables in the databases 316. Such automatic processing and methodology reports are performed and generated by the enterprise server 314, and is described in detail below.

In step 9516, operators 4514 analyze the customer provided data 4504 in the first part 4510 of the databases 4508. The operators 4514 may also analyze the other information 4516. In performing this analysis, the operators 4514 may refer to the methodology reports 4518 run in step 9514. Since these methodology reports 4518 were prepared by the enterprise server 314, the system of the invention assists the operators 4514 in performing this analysis. Based on the

-114-

analysis of the operators 4514, additional database information for the databases 4508 is generated.

In step 9518, this additional database information is loaded into the second part 4512 of the databases 4508. The databases 4508, at that point, are fully loaded. Periodically, the steps of flowchart 9502 must be repeated in order to update the databases 4508 with additional and/or modified customer provided data 4504 and/or additional and/or modified other information 4516.

In an alternate embodiment of the invention, data is not preloaded into the invention's databases. Instead, the invention accesses the customer's corporate databases for data on an as needed basis. This alternate embodiment is described below with respect to the BOM databases 626, but are applicable to the other tables in the databases 316 as well.

Document Bibliographic Databases

FIGS. 12B-12F and 12H illustrate the structure of the document bibliographic databases 602. As indicated in FIG. 6, the document bibliographic databases 602 include bibliographic databases for documents of interest to the customer. The patent bibliographic databases 604 include information about (i.e., bibliographic information) U.S. and/or foreign patents. Preferably, the patent bibliographic databases 604 have bibliographic information on all U.S. Patents and a subset of all foreign patents. As an alternative embodiment, the patent bibliographic databases 604 include (in addition to information on foreign patents) patent bibliographic information on only a subset of all U.S. patents, such as all U.S. patents available in electronic form from the U.S. Patent and Trademark Office, or all U.S. patents that issued after a certain date, or all U.S. patents of interest to the customer.

The patent bibliographic databases 604 include a patent table 1222 (FIG. 12H). The patent table 1222 includes a record for each U.S. and foreign patent

represented in the patent bibliographic databases 604. Each record in the patent database 1222 includes a document_id attribute that stores a unique identifier (or key) for the associated patent. It is noted that, in the tables of the databases 316, the symbol FK stands for foreign key, AK stands for alternate key, and IE stands for inversion entry (which is a non-unique index).

Each record of the patent database 1222 also includes attributes that, for the most part, correspond to the bibliographic information on the first page of U.S. patents. In an embodiment of the invention, each record of the patent database 1222 includes attributes that, for the most part, correspond to the bibliographic information contained in the electronic representations of U.S. patents publicly available from the U.S. Patent and Trademark Office.

For example, a record in the patent database 1222 includes a document_number attribute that stores a patent number 4004 (see the example patent in FIG. 40). In an entry of the patent database 1222, the AppNo attribute corresponds to the application number 4014, the AppDate corresponds to the filing data 4016, the title corresponds to the title of the patent 4010, the issue date corresponds to the date that the patent issued 4006, the NumClaims corresponds to the number of claims in the patent 4036, the AsstExaminerLastName corresponds the last name of the assistant examiner 4032, the AsstExaminerFirstName corresponds to the first name of the assistant examiner 4032, the PrimaryExaminerLastName corresponds to the last name of the primary examiner 4030, and the PrimaryExaminerFirstName corresponds to the first name of the primary examiner 4030.

Also in each record of the patent database 1222, the NumDrawingPages corresponds to the number of drawing sheets 4038, the disclaimer date corresponds to any terminal disclaimer 4106 (FIG. 41), the ReissueLevel, ReissueAppNo, ReissueAppDate, ReissuePatentNo, and ReissueIssueDate corresponds to any reissue information 4308, 4304, and 4306 (FIG. 43).

-116-

Each record of the patent database 1222 also includes attributes that correspond to patent bibliographic information not shown on the front page of U.S. patents. For example, a record of the patent database 1222 also includes a SeriesCode that corresponds to the series code of the patent. Other information
5 contained in each record of the patent database 1222 and not shown on the front page of the U.S. patent is the AppType, PubLevel, ArtUnit, ExemplaryClaimNo, NumFigures, NumSpecPages, TermYears, and IntlEdition. Each record of the patent database 1222 may also include fields whose values are calculated during the loading phase. For example, each record of the patent database 1222 may
10 include a calc_exp_date that corresponds to the expiration date of the patent. This date is calculated and loaded into the patent database 1222 during the load phase of the patent database 1222 (described below). calc_exp_date and issue date are collectively referred to as patent term expiration related information.

Each record of the patent database 1222 also includes one or more
15 user_defined fields. Users may enter any information into this field. The amount of information that can be entered into this field is relatively large, such as 32 kbytes or greater. This field is preferably indexed searchable. The user can enter into this field information that is specific and/or of interest to his company. For example, a user may enter into this field its own matter or reference/tracking
20 number.

Additionally, the invention allows operators to add any number of additional user defined fields, both into the patent database 1222 and into any other table of the databases 316. The fields must be of certain predefined types, such as date fields, string fields, numeric fields, etc. The user can define the name
25 of these fields and the types of these fields (from a number of available field types). Preferably, these fields are indexed and searchable.

The record in the patent table 1222 corresponding to a particular patent is herein called the base record for the patent (because this record in the patent table 1222 includes most of the bibliographical information about the patent). The

patent bibliographic databases 604 include other tables that store additional patent bibliographic information about each patent represented in the patent bibliographic databases 604. Records in these other tables are linked to their respective base records in the patent database 1222 via the document_id attribute.

5 An assignee table 1201 (FIG. 12B) includes information on the assignees of a patent, if any. A given patent may have multiple assignees. For each assignee of a patent, there is a record in the assignee table 1201. These assignee records in the assignee table 1201 are linked to the corresponding base record in the patent table 1222 via the document_id attribute. Each record of the assignee
10 table 1201 includes an assignee_id attribute which is an identifier that uniquely identifies the assignee. Each record of the assignee table 1201 also includes information pertaining to the assignee, such as country information, state information, the name of the assignee, and the city and zip code of the assignee. This information is found on the front page of U.S. patents (see field 4104 in FIG.
15 41). In each entry of the assignee table 1201, the country and state of the assignee as preferably specified using codes. These codes are defined in a state table 1207 and a country table 1210 (FIG. 12C).

 An intlclass table 1203 (FIG. 12B) stores information pertaining to the international class of a patent and the international search classes of the patent.
20 For a given patent, the intlclass table 1203 includes a record for each international class to which the patent is assigned. Additionally, the intlclass table 1203 includes a record for each international search class which was searched during the prosecution of the patent. Whether a record in the intlclass table 1203 corresponds to an international class or an international search class is denoted by
25 the attribute is_search_class. This attribute is set to true if the record corresponds to an international search class. Referring to FIG. 40, the international class is identified by reference number 4024. Not all patents have international search classes, and this is the case with the patent shown in FIG. 40. The records in the

-118-

intlclass table 1203 are linked to the associated base record in the patent table 1222 via the document_id attribute.

The patent_class_xref table 1204 (FIG. 12B) includes information on the U.S. classification of a patent. For a given patent, the patent_class_xref table 1204 includes an entry for the original classification of a patent. The patent_class_xref table 1204 also includes an entry for each unofficial classification of the patent, and each digest classification of the patent. Whether or not a record in the patent_class_xref table 1204 corresponds to the original classification is denoted by an original attribute. Whether or not an entry in the patent_class_xref table 1204 corresponds to an unofficial classification or a digest classification is designated by a patent_class_type_id code, whose values are defined by a patent_class_type table 1202. Records in the patent_class_xref table 1204 are linked to the corresponding base record in the patent table 1222 via the document_id attribute.

A patent class/subclass is stored in a record of the patent_class_xref table 1204 using the patent_class_id attribute, the subclass_id attribute, and the suffix_id attribute. Consider the following class/subclass: 364/419.19. For this example, the patent_class_id is equal to 364. The subclass_id is equal to 419. The suffix_id attribute is equal to 19. By breaking the class/subclass into these three fields, it is possible to fine tune searches and direct searches to any combination of the class, subclass, or subclass suffix.

The patent_class_id attribute is actually a code. The same is true of the subclass_id and the suffix_id attributes. These patent class codes are defined in a patent_class table 1209 (FIG. 12C).

A RelatedApp table 1206 (FIG. 12B) stores information on applications which are related to a patent. For a given patent, the RelatedApp table 1206 includes a record for each application that is related to the patent. Related application data is shown, for example, in FIG. 44B at reference number 4490. Records in the RelatedApp table 1206 are linked to the associated base record in

the patent table 1222 via the document_id attribute. An entry in the RelatedApp table 1206 includes attributes to store the serial number of the related application, the filing date of the related application, the status of the related application, and the patent number and issue date of the related application, if the related application issued as a patent. Each entry of the RelatedApp table 1206 also includes a grammarcode attribute that stores a code corresponding to such text as "continuation of", "continuation-in-part," "which is a continuation-in-part of". These grammar codes are found in the electronic representations of U.S. patents publicly available from the U.S. Patent and Trademark Office.

A LegalRepAttor table 1205 includes information on the attorney or agent who prosecuted the patent. Such information is shown in FIG. 40 at reference number 4034. There is a record in the LegalRepAttor table 1205 for each attorney or agent who prosecuted and is listed on the front page of the patent. Records in the LegalRepAttor table 1205 are linked to the corresponding base record in the patent table 1222 via the document_id attribute.

Referring to FIG. 12C, a PatentRef table 1208 stores information on U.S. patents that were cited during the prosecution of a given patent. The PatentRef table 1208 includes a record for each U.S. patent that was cited during the prosecution of a given patent. Such reference to U.S. patents are shown, for example, in FIG. 40 at reference number 4028. Each record of the PatentRef table 1208 includes a RefPatentNo attribute that represents the patent number of the reference patent. Each record of the PatentRef table 1208 also includes attributes that store the issue date of the reference patent, the first named inventor of the reference patent, and the class/subclass of the reference patent. Records in the PatentRef table 1208 are linked to the corresponding base record in the patent table 1222 using the document_id attribute.

Referring now to FIG. 12D, a SearchClass table 1211 stores information on U.S. classes and subclasses which were searched during the prosecution of a patent. For any given patent, the SearchClass table 1211 includes a record for

-120-

each class/subclass that was searched during the prosecution of the patent. U.S. search class information is shown, for example, in FIG. 40 at reference number 4026. Records in the SearchClass table 1211 are linked to the corresponding base record in the patent table 1222 via the document_id attribute.

5 The inventor table 1212 includes information on the inventors of a patent. The inventor table 1212 includes a record for each inventor of a given patent. Records in the inventor table 1212 are linked to the corresponding base record in the patent table 1222 via the document_id attribute. Inventorship information is shown, for example, in FIG. 44A by reference number 4450. Each record in the
10 inventorship table 1212 includes an inventor_id attribute that stores a key that uniquely identifies the inventor. Each record of the inventorship table 1212 also includes attributes to identify the first and last name of the inventor, the address of the inventor, and the state and country of the inventor. The state and country values are specified by the state_id and the country_id attributes, which are codes.

15 The state and country codes are defined in the state table 1207 and the country table 1210, respectively (FIG. 12C).

Referring now to FIG. 12E, the LegalRepFirm table 1213 includes information on the law firm that prosecuted the patent. There is one record in the LegalRepFirm table 1213 for each law firm that prosecuted the patent and that is
20 shown on the front page of the patent. This law firm information is shown, for example, in FIG. 40 by reference number 4034. Records in the LegalRepFirm table 1213 are linked to the corresponding base record in the patent table 1222 via the document_id attribute.

25 The PatCoopTreaty table 1214 stores information on a PCT application, in those cases where the U.S. patent was first filed as a PCT application. Such PCT information is shown, for example, in FIG. 42 at reference number 4204. Records in the PatCoopTreaty table 1214 are linked to the corresponding base record in the patent table 1222 via the document_id attribute.

-121-

The priority table 1215 includes priority information related to the patent. Such priority information is shown, for example, in FIG. 43 at reference number 4310. Records in the priority table 1215 are linked to the corresponding base record in the patent table 1222 by using the document_id attribute.

5 Referring now to FIG. 12F, the ForeignRef table 1216 includes information on foreign references that were cited during the prosecution of the patent. Such citations of foreign references are shown, for example, in FIG. 40 at reference number 4040. For a given patent, the ForeignRef table 1216 includes a record for each foreign reference that was cited against the patent. Each record
10 of the ForeignRef table 1216 includes a ForeignPatentNo attribute that corresponds to the patent number of the foreign reference. Each record of the ForeignRef table 1216 also includes attributes that store information identifying the issue date of the foreign patent, the country of the foreign patent, and U.S. and international classes.

15 Referring again to FIG. 6, the document bibliographic databases 602 include other bibliographic databases corresponding to documents that are of interest to the customer. In the example of FIG. 6, the document bibliographic databases 602 include deposition bibliographic databases 606, pleadings bibliographic databases 608, and prior art bibliographic databases 610. These
20 databases may be of interest to the customer if, for example, the customer is involved in a patent litigation. Examples of the database structure of the deposition bibliographic databases 606, the pleadings bibliographic databases 608, and the prior art bibliographic databases 610 are shown in FIG. 12H. A depositions table 1225 corresponds to the deposition bibliographic databases 606.
25 The depositions table 1225 includes a record for each deposition-related document that is of interest to the customer. Each record in the depositions table 1225 includes a document_id attribute that is a key that uniquely identifies the document. Each record of the depositions table 1225 also includes other bibliographic information about the deposition-related document, such as a

-122-

5 deposition_id attribute that stores an identifier that identifies the deposition document. Each record of the depositions table 1225 may also include other bibliographic information such as the date that the deposition took place, the parties involved in the deposition, the number of pages of the deposition transcript, the place that the deposition took place, a summary of the deposition, etc.

10 A pleadings table 1224 corresponds to the pleadings bibliographic databases 608. The pleadings table 1224 includes a record for each pleadings-related document that is of interest to the customer. Each record of the pleadings table 1224 includes a document_id attribute that stores a key that uniquely identifies the pleadings-related document. Each record of the pleadings database 1224 may also include other bibliographic information, such as a pleadings_id attribute. The pleadings table 1224 may include other attributes, such as attributes that identify the author of the pleadings, the date that the pleadings were created, 15 the subject matter of the pleadings, the number of pages of the pleadings, the parties involved in the pleadings, the title of the pleadings, the type of the pleadings, etc.

The prior_art table 1226 corresponds to the prior art bibliographic 20 databases 610. There is a record in the prior_art table 1226 for each prior art reference that is of interest to the customer. Each reference of the prior_art table 1226 includes a document_id attribute that uniquely identifies the prior art reference. Each record of the prior_art table 1226 also includes other attributes, such as a prior_art_id attribute, and attributes to store information identifying the author of the prior art reference, the publication date of the prior art reference, the 25 title of the prior art reference, the subject matter of the prior art reference, the number of pages of the prior art reference, etc.

Still referring to FIG. 12H, a document table 1223 represents a table of the methodology support databases 642 (FIG. 6). The document table 1223 represents a master index of all of the documents represented in the document

5 bibliographic databases 602. The document table 1223 has a record for each document represented in the patent bibliographic databases 604, the deposition bibliographic databases 606, the pleadings bibliographic databases 608, the prior art bibliographic databases 610, and any other bibliographic databases in the document bibliographic databases 602. Each record of the document table 1223 includes a document_id attribute that is a key that uniquely identifies a document. The document_id attribute is a common attribute in the document bibliographic databases 602, and serves to link the tables in the document bibliographic databases 602 together. Each record of the document table 1223 also includes a document_type attribute that identifies the type of the document (for example, whether the document is a patent, a deposition-related document, a pleadings-related document, a prior art-related document, etc.). Each record of the document table 1223 also includes a document_guid attribute that serves as a second identifier for the document, and a version_number attribute that identifies the version of the document.

15 The document bibliographic databases 602 shall now be further described with reference to an example shown in FIGS. 13-17. FIG. 13 illustrates an example patent table 1222X. For simplicity, the example patent table 1222X only shows the document_id attribute and the document_number. The patent table 1222X stores therein bibliographic information on documents D1-D14, each of which corresponds to a patent.

20 FIG. 14 is an example depositions table 1225X. For simplicity, only the document_id and the deposition_id attributes are shown. The depositions table 1225X stores bibliographic information on documents D15-D17, each of which represents a deposition-related document.

25 FIG. 15 shows an example pleadings table 1224X. For simplicity, only the document_id and the pleadings_id attributes are shown. The pleadings table 1224X stores information on documents D18-D21, each of which corresponds to a pleadings-related document.

-124-

FIG. 16 is an example prior_art table 1226X. For simplicity, only the document_id and the prior_art_id attributes are shown. The prior_art table 1226X stores information on documents D22 and D23, each of which represents a prior art-related document.

5 FIG. 17 is an example document table 1223X. The document table 1223X includes an entry for each of the documents represented in the example patent table 1222X, the example depositions table 1225X, the example pleadings table 1224X, and the example prior_art table 1226X (shown in FIGS. 13-16, respectively). Accordingly, the document table 1223X has entries for documents
10 D1-D23. For simplicity, only the document_id and the document_type attributes are shown in the example document table 1223X. The document_type attribute indicates the type of the document.

The preferred procedure for loading the document bibliographic databases 602 shall now be described.

15 FIG. 46 illustrates a data flow diagram of the loading of the patent bibliographic databases 604. FIG. 96 is a flowchart 9602 that corresponds to the data flow diagram 4602 of FIG. 46, and that represents the operation of the invention when loading bibliographic information of a given patent into the patent bibliographic databases 604.

20 In step 9606, the patent text file 4604 and the patent image file 4624 of the patent being processed is received. For U.S. patents, electronic text and image files of patents are available from the U.S. Patent and Trademark Office. Similarly, text and image files corresponding to foreign patents are available from foreign patent offices.

25 In step 9608, an error detection and correction module 4612 in a parser 4610 analyzes the patent text file 4604 and detects any errors in the bibliographic data contained in the patent text file 4604. The error detection and correction module 4612 detects many types of errors. For example, an error which is detected by the error detection and correction module 4612 involves the issue

-125-

date. With regard to patent text files 4604 received from the U.S. Patent Office, the issue date appears in the following format:

ISD 19790315.

Thus, the issue date information is set off by the ISD token. This token is followed by four digits representing the year, two digits representing the month, and two digits representing the day. The error detection and correction module 4612 detects errors in the date given known date ranges. For example, the error detection and correction module would identify the following as an error:

ISD 99790315

The error detection and correction module 4612 would detect this as an error, since the year value is out of range. The error detection and correction module 4612 operates in this manner to detect all date-related errors.

Another error which the error detection and correction module 4612 detects involves state information and zip code information. The error detection and correction module 4612 correlates the state information and zip code information contained in the patent text file 4604 to determine if the zip code is possibly incorrect given the state, and vice versa. Such processing is possible, since each state is associated with a finite set of zip codes.

Another error which the error detection and correction module 4612 detects involves the patent number. The patent number of a U.S. patent comprises seven digits. The error detection and correction module 4612 confirms that the patent number has seven digits, and detects an error if it does not. The error detection and correction module 4612 performs a similar function for non-U.S. patents, and other types of patent documents, such as reissues, reexams, etc.

The patent number of a reissued patent has the token RE followed by five digits. The error detection and correction module 4612 detects an error if the patent number of a reissued patent does not follow this format.

The detection performed by the error detection and correction module 4612 also involves street addresses. The error detection and correction module

4612 detects a possible error if the street address has no alphabetic characters, or has no space characters, or has no digit characters. Similarly, the error detection and correction module 4612 detects a possible error with the city name if the information in the patent text file 4604 corresponding to a city name has no alphabetic characters.

The error detection and correction module 4612 performs additional error correction with regard to U.S. states by comparing information corresponding to U.S. state names from the patent text file 4604 with a list of the U.S. state abbreviations used by the U.S. Patent Office. If the state information does not match one of the state abbreviations, the error detection and correction module 4612 detects an error.

Similarly, the error detection and correction module 4612 detects errors with information in the patent text file 4604 corresponding to country abbreviations by comparing such country information from the patent text file 4604 with the list of country codes utilized by the U.S. Patent and Trademark Office (and any other appropriate foreign patent office).

The error detection and correction module 4612 also checks to ensure that mandatory data is contained in the patent text file 4604. Some bibliographic fields in the patent text file 4604 are mandatory, while others are optional. For example, the patent text file 4604 must contain document type information that identifies whether the patent is a utility patent, a design patent, a plant patent, a defensive publication, etc. In contrast, the patent text file 4604 may or may not include assignee information, since a patent does not have to have an assignee. The error detection and correction module 4612 checks for mandatory information, and detects an error if mandatory information is not contained in the patent text file 4604.

Many fields in the patent text file 4604 must be in a proper format. Some of these fields were already discussed above, such as the patent number and date. Other fields which must be in a certain format are classes and subclasses, and

-127-

dates. The error detection and correction module 4612 determines whether this information from the patent text file 4604 is in the proper format, and detects an error if they are not.

Also in step 9608, the error detection and correction module 4612 automatically corrects the errors that it detects, where such automatic correction is possible. For example, given the patent number, the error detection and correction module 4612 can estimate the issue date. If the error detection and correction module 4612 detected an error with the issue date, then the error detection and correction module 4612 can attempt to automatically correct the error in the issue date based on the patent number (assuming that the patent number was found to contain no errors).

If the error detection and correction module 4612 cannot automatically correct the errors that it detected, then the error detection and correction module 4612 notifies an operator 4616. The operator 4616 then manually corrects the error.

The error detection and correction module 4612 also detects for errors in the information in the patent text file 4604 corresponding to assignees. The same company can be listed as assignee in different patents using different names. For example, IBM can be listed as the assignee in patents using various different names, such as IBM; International Business Machines; IBM, Inc.; etc. The error detection and correction module 4612 includes a thesaurus feature that lists the most commonly used names for well-known companies. The error detection and correction module 4612 compares the assignee name from the patent text file 4604 with this thesaurus and replaces the assignee information in the patent text file 4604 with the name retrieved from the thesaurus. Alternatively, the name retrieved from the thesaurus may be written to an appropriate user-defined field of the bibliographic databases 602.

In step 9610, a formatting module 4614 in the parser 4610 formats the error detected and corrected patent text file 4604. This formatting is performed

-128-

in order to convert the format of the patent text file 4604 to the format of the patent bibliographic databases 604. Such formatting includes formatting performed with respect to patent numbers. The required formats for different types of patents are presented in Table 4, below. It is noted that Table 4 is not a complete list, but just an illustrative one. Other types of patents (applicable in both the U.S. and foreign countries) could also be processed by the formatting module 4614.

Table 4

Type of Patent	Format
Utility patent	7
Design patent	D6
Plant patent	PP5
Utility SIR	H7
Design SIR	HD6
Plant SIR	HP6
Defensive publication	T6
X patent	X7
Additional improvements patent	AI4
Non-patent literature	NP6
Reissued utility patent	RE5
Reissued design patent	RD5
Reissued plant patent	RP5
Reissued X patent	RX5

The formatting module 4614 confirms that the information in the patent text file 4604 corresponding to patent numbers corresponds to the formats shown in Table 4. If the format is not as shown in Table 4, the formatting module 4614

modifies the patent text file 4604 so that the patent number conforms with the format of Table 4. For example, consider the format of a plant patent, which is a token PP immediately followed by five digits. Assume that the information in the patent text file 4604 was as follows: PP-11111 (for illustrative purposes, spaces are indicated as dashes). This patent number is not in the proper format because five digits do not immediately follow the PP token. Instead, there is a space between the PP token and the five digits. Accordingly, the formatting module 4614 modifies the patent text file 4604 so that the patent number is in the proper format. In the example above, the formatting module 4614 deletes the space between the PP token and the five digits.

In step 9612, the parser 4610 generates a corrected and normalized patent text file 4618 from the error detected and corrected, and formatted patent text file 4604. The corrected and normalized patent text file includes both the text from the patent text file 4604 and the error detected and corrected, and formatted patent bibliographic information from the patent text file 4604.

When generating the corrected and normalized patent text file 4618, the parser 4610 may take into consideration other patent information 4605. The parser 4610 may modify the information from the patent text file 4604 with this other patent information 4605 when generating the corrected and normalized patent text file 4619. For example, a patent is sometimes reassigned to another company after it issues. The new assignee, however, is not indicated on the front page of the patent (because the patent was published before the reassignment took place). Often, patent reassignment information is in the assignment records of the U.S. Patent and Trademark Office (if such information was recorded with the USPTO). Such assignment records may constitute part of the other patent information 4605. The parser 4610 in generating the corrected and normalized patent text file 4618 may take into account such assignment information when generating the corrected and normalized patent text file. For example, the parser 4610 may include a note in the corresponding record of the corrected and normalized patent text file 4618

-130-

that the patent has been reassigned, and indicate the new assignee(s). This note may be, for example, in a field of the corrected and normalized patent text file 4618 corresponding to the user defined attribute in the record of the patent table 1222 (FIG. 12H).

5 The patent bibliographic information in the corrected and normalized patent text file 4618 is stored in normalized format. Preferably, this normalized format is a condensed, field delimited format. Essentially, in normalizing the patent bibliographic information from the patent text file 4604, the parser 4610 separates the patent bibliographic information into fields that correspond to the
10 attributes in the patent bibliographic databases 604. These fields are delimited either using commas, tabs, or some other symbol.

 For example, the format of a class/subclass in the patent text file 4604 is as follows: XCL---17-34-2 (for illustrative purposes, spaces are indicated as dashes). Therefore, in the patent text file 4604, a class/subclass is denoted by the
15 token XCL, followed by two spaces. Three characters follow this token that correspond to the class. This is followed by three characters that correspond to the subclass. Finally, two characters follow the subclass. These two characters correspond to the suffix. In the corrected and normalized patent text file 4618, this class/subclass information is represented in a field delimited format, as follows: 17,
20 34, 2. This is a comma delimited format. In other embodiments, the corrected and normalized patent text file 4618 uses a tab delimited format for the patent bibliographic information.

 In step 9614, the corrected and normalized patent text file 4618 is used to regenerate the patent text file 4604. The regenerated patent text file is designated
25 by reference number 4608. A comparison module 4606 then compares the patent text file 4604 with the regenerated patent text file 4608. If the regenerated patent text file 4608 differs from the patent text file 4604, then the comparison module 4606 determines that the parser 4610 incorrectly generated the corrected and normalized patent text file 4618 from the patent text file 4604. If such an error is

-131-

detected, then the comparison module 4606 notifies an operator 4616. It is noted that step 9614 is optional.

Preferably, the corrected and normalized patent text file 4618 represents an object-oriented framework. Generating the corrected and normalized patent text file 4618 as a framework is useful because it can then be easily exported to a number of different formats, for later processing. Accordingly, in step 9616, the corrected and normalized patent text file is stored in various formats for later retrieval. Such formats may include a human editable production format 4620, and/or a SPML format 4622. It is noted that step 9616 is optional.

In step 9618, the corrected and normalized patent text file is paginated by a pagination module 4626 to generate a patent equivalent text file 4630. The patent equivalent text file 4630 includes equivalency information that establishes an equivalency relationship between the text in the patent equivalent text file 4630 and the image in the patent image file 4624. For example, this equivalency information includes pagination information that enables the patent equivalent text file 4630 to be displayed having the same pagination (line breaks, column breaks, page breaks) as the patent image file 4624. The pagination module 4626 generates the patent equivalent text file 4630 by comparing the patent text in the corrected and normalized patent text file 4618 with the patent image file 4624 to detect equivalency information. This equivalency information is then embedded in the patent equivalent text file 4630, along with the patent text. While the pagination module 4626 is capable of performing the pagination operation automatically, in some cases some manual intervention is required. In accordance, an operator 4628 is sometimes involved with the pagination process performed by the pagination module 4626. The pagination process performed by the pagination module 4626 is further described in U.S. Patent No. 5,623,681, U.S. Patent No. 5,623,679, and pending U.S. Patent Application Serial No. 08/341,129, all of which are incorporated by reference herein. It is noted that step 9618 is optional.

-132-

In step 9620, a patent bibliographic flat file 4619 is generated from the corrected and normalized patent text file 4618. The patent bibliographic flat file 4619 includes the normalized patent bibliographic information from the corrected and normalized patent text file 4618. Preferably, the patent bibliographic flat file 4619 does not include the patent text. The patent bibliographic flat file 4619 is generated using standard database utilities.

In step 9622, the patent bibliographic information in the patent bibliographic flat file 4619 is loaded into the tables of the patent bibliographic databases 604. This loading operation is performed by using standard database loading utilities.

After completion of step 9622, the patent bibliographic information from the patent being processed is completely loaded into the patent bibliographic databases 604. The steps of flowchart 9602 are repeated for each patent that is to be represented in the patent bibliographic databases 604.

Referring again to FIG. 6, the loading of the deposition bibliographic databases 606, the pleadings bibliographic databases 608, and the prior art bibliographic databases 610 (and any other bibliographic databases in the document bibliographic databases 602) is conducted generally according to the methodology shown in FIGS. 46 and 96 for U.S. patents. More generally, the extract and load for these bibliographic databases may be performed according to the methodology shown in FIGS. 45 and 95, described above. The actual implementation of the extract and load methodologies for these bibliographic databases will necessarily be case specific and depend on a number of factors, including the bibliographic information contained in the source documents, the format of the source documents (that is, whether or not they are electronically stored, the storage format, how bibliographic data is represented in the storage format, etc.), etc. The extract and load processes for these other bibliographic databases will be apparent to persons skilled in the relevant art(s) based on the discussion contained herein.

Group Databases

The group databases 621 are described below. In particular, the structure of the group databases 621 and the methodologies for loading the group databases 621 with data are described in the following sections.

5 *User Defined Groups*

FIG. 12I depicts the user-defined group (also called arbitrary group) databases 624. A group_table database 1227 has a record for each user-defined group. Each record in the group_table database 1227 includes a group_id attribute that stores a key that uniquely identifies the user-defined group from all other
10 user-defined groups. Each record of the group_table database 1227 also includes a group_guid which is an alternate key (identifier) of the record, a name attribute that stores the name of the user-defined group, a creation_date attribute that stores the date that the group was created, a modification_date attribute that stores the date that the group was last modified, a version_number attribute that
15 stores the version number of the group, a description attribute that stores a description of the group, and an owner_user_id attribute that identifies the owner (that is, the creator) of the user-defined group.

According to the present invention, groups are hierarchically organized. In other words, a given group can be a child of one or more parent groups, and
20 can also be a parent to one or more child groups. This hierarchical organization is illustrated by way of example in FIG. 18, which is an example screen shot 1801 displayed by the client grouping module 712 on the client display unit 1122. This screen shot 1801 includes a first window 1802 and a second window 1804. The first window 1802 graphically depicts the hierarchical organization of the groups,
25 both user-defined and predefined. The highest level group is called the root group,

or repository. A BOM group is a child of the root group. The BOM group is a pre-defined group and is discussed below.

A user-defined group is also a child of the root group. The user-defined group has a number of child groups. These child groups are user-defined groups. In the example, the child groups are: companies to acquire, patents about bikes, and litigation. The companies to acquire user-defined group has a number of child user-defined groups. They include: ABC Corp., XYZ Corp., and PQR Corp.

The hierarchical structure of the user-defined groups is defined by information in a group_group_xref table 1229 (FIG. 12I). The group_group_xref table 1229 includes a record for each parent/child relationship in the user-defined group hierarchical structure. Each record of the group_group_xref table 1229 includes a parent attribute and a group_id attribute. Both the parent attribute and the group_id attribute store group_id values. The parent attribute stores the group_id of the parent group, and the group_id stores the group_id of the child group.

The group_table database 1227 and the group_group_xref table 1229 shall be further described with reference to an example illustrated in FIGS. 19 and 20. FIG. 19 illustrates an example group_table database 1227X that includes information representative of the user-defined group hierarchical structure shown in FIG. 18. The example group_table database 1227X in FIG. 19 has a record for each user-defined group shown in FIG. 18. Accordingly, there is an entry in the group_table database 1227X for the following user-defined groups: companies to acquire, ABC Corp., XYZ Corp., PQR Corp., patents about bikes, and litigation. The group_id values of these user-defined groups are GRP1-GRP6, respectively. Note that the group_table database 1227X has no information pertaining to the hierarchical structure of the user-defined groups. Instead, such information on the hierarchical structure of the user-defined groups is contained in the group_group_xref table 1229.

-135-

FIG. 20 illustrates an example group_group_xref table 1229X. The group_group_xref table 1229X has entries that define parent/child relationships in the user-defined groups hierarchical structure. For example, an entry 2002 specifies that the parent attribute is GRP1 (corresponding to the companies to acquire group), and the group_id attribute is GRP2 (corresponding to the ABC Corp. group). Therefore, entry 2002 in the group_group_xref table 1229X specifies that the ABC Corp. group is a child of the companies to acquire group. Similarly, the entry 2004 specifies that the XYZ Corp. group is a child of the companies to acquire group, and entry 2006 specifies that the PQR Corp. group is a child of the companies to acquire group.

Referring again to FIG. 12I, a user-defined group can contain any number of documents, such as patents. The documents in a group are considered to be relevant to the group. In practice, a user places the documents in a group.

A group_document_xref table 1228 stores information that identifies the documents that are in each user-defined group. The group_document_xref table 1228 includes a record for each document in each user-defined group. Each record of the group_document_xref table 1228 stores a group_id attribute to identify the user-defined group, a document_id attribute to identify a document in that group, and a modification_date attribute that stores the date that the record in the group_document_xref table 1228 was last modified.

The group_document_xref table 1228 shall be further described with reference to an example shown in FIG. 21. FIG. 21 illustrates an example group_document_xref table 1228X. The example in FIG. 21 indicates that documents D1-D3 are in the user-defined group having the group_id of GRP2 (that is, the ABC Corp. group). Documents D7 and D8 are in the XYZ Corp. group; documents D11 and 12 are in the PQR Corp. group; documents D1-D3 and D9 are in the patents about bikes group; documents D22, D23, D18, D21, D16 and D1 are in the litigation group.

Referring again to FIG. 18, note that the user has selected the ABC Corp. group in the window 1802 (this is indicated by the dotted circle 1850). Selection of a group in window 1802 results in a listing of the documents in the selected group being displayed in a window 1804. Accordingly, a list including documents D1-D3 is displayed in window 1804. Other bibliographic information pertaining to these documents can also be displayed in window 1804 in a tabular or "spreadsheet" format by retrieving such bibliographic information from the appropriate tables in the document bibliographic databases 602.

A user can view a document by double-clicking (or use any other well known GUI technique) on that document in the window 1804. In the example of FIG. 18, the user has selected document D1 (indicated by dotted circle 1852). This results in the document being displayed in a window 1806. The window 1806 includes a window 1808, where the text of document D1 is displayed, and/or a window 1810, where the image of document D1 is displayed. The example of window 1806 where text and images of a document are selectively displayed is more particularly shown in FIG. 112. An example of screen shot 1801 where the user-defined group hierarchical structure is shown in one window 1802 and a list of documents is displayed in another window 1804 is more particularly shown in FIG. 58.

The procedures for initially loading the user-defined group databases 624 with data generally track the generic extract and load procedures illustrated in FIGS. 45 and 96. Such an initial extract and load procedure will only be required if any user-defined groups exist prior to initiation of the enterprise server 314 for a particular customer. If the customer does not have any user-defined groups prior to initiating the enterprise server 314 for the first time, then the user-defined group databases 624 will be loaded with information as the user creates and modifies user-defined groups.

Predefined Group Databases

The predefined group databases 622 (also called system defined group databases) are described in the following sections.

Bill of Materials (BOM) Databases

5 A bill of materials (BOM) is a well-known data structure often used by companies to abstractly represent a product. A BOM is a hierarchical and recursive data structure that identifies the subassemblies of a product, and that identifies the parts of the subassemblies. Accordingly, a product's BOM is useful for recording the sub-assemblies and parts needed to construct the product.
10 BOMs are also called herein BOM data structures, or BOM structures, or BOM trees.

FIG. 22 illustrates an example BOM 2202 for a bicycle. The bicycle is the product (also called the assembly). As indicated in FIG. 22, the bicycle includes a frame, two wheels, and a handle bar (only a portion of the BOM for the bicycle is shown for simplicity purposes).
15

The frame, the two wheels, and the handle bar are subassemblies because they each are composed of one or more subassemblies or parts. For example, the frame subassembly includes a screw, which is a part. A part is defined as an item which does not have component parts separately represented in the BOM. Each
20 of the two wheel subassemblies includes a rim and a tire, which are both parts. The handlebar subassembly includes a metal rod and two hand grips, all three of which are parts.

Although not indicated in the example of FIG. 22, a subassembly may be composed of any combination of parts and/or subassemblies.

25 The BOM 2202 in FIG. 22 includes a number of nodes, called BOM nodes. A node corresponds to each assembly (i.e., the bicycle node), each

-138-

subassembly (i.e., the frame, wheel, and handle bar nodes), and each part (i.e., the screw, rim, tire, metal rod, and hand grip nodes).

In the context of and as applied by the present invention, each BOM node has a logical level in the hierarchical structure of the BOM. According to an embodiment of the invention, there are three BOM logical levels (although a customer can define others): assembly, subassembly, and part. Accordingly, the logical level of the bicycle node is assembly. The logical level of the frame, wheel, and handle bar nodes is subassembly. The logical level of the screw, rim, tire, metal rod, and hand grip is part.

In the invention, the reason for using logical levels is to enable enhanced reporting and analysis functions, such as the reporting functions described herein that are organized according to BOM logical level. As discussed above, an advantage of predefined groups is that their structure and attributes are better defined and more extensive than user-defined groups. Accordingly, more extensive and more useful analysis and reporting functions can be developed for predefined groups.

FIG. 23A illustrates an example BOM 2302 of a lawn mower. The lawn mower includes a frame, an engine, and a handlebar. These are all subassemblies (although for simplicity purposes the engine subassembly is not further defined in the BOM 2302 shown in FIG. 23A). The frame subassembly includes a screw. The handlebar subassembly includes two hand grips. The screw and the hand grips are parts. It is noted that the screw part in the bicycle BOM 2202 is the same physical part as the screw part in the lawnmower BOM 2302. Similarly, the hand grip part in the bicycle BOM 2202 is the same physical part as the handgrip part in the lawnmower BOM 2302. However, the frame subassembly in the bicycle BOM 2202 is not the same physical part as the frame subassembly in the lawnmower BOM 2302. Also, the handle bar subassemblies in the bicycle BOM 2202 and the lawnmower BOM 2302 are not the same.

FIG. 23B illustrates a portion of a BOM of an example antique car. This BOM includes a body subassembly, and two wheel subassemblies. In the present example, these wheel subassemblies are the same as the wheel subassemblies in FIG. 22. As discussed below, these wheel subassemblies are represented as a single BOM group in the BOM databases 626. Both the BOM group corresponding to the bicycle and the BOM group corresponding to the antique car include the wheel BOM group (that is, the wheel BOM group is a child group of the bicycle BOM group and the antique car BOM group). That is, the wheel BOM group is shared by the bicycle BOM group and the antique car BOM group. Accordingly, the wheel BOM group is herein called a shared group.

The BOM databases 626 are illustrated in FIG. 12G. A BOM table 1217 includes a record for each unique BOM node of each BOM of interest. Accordingly, each record in the BOM table 1217 represents a BOM node in a BOM of interest to the customer.

The BOM table 1217 includes entries for only unique BOM nodes of interest. Again referring to FIG. 22, the two wheel subassembly BOM nodes are not unique with respect to each other. Therefore, the BOM table 1217 only includes a single record for the wheel BOM node. Similarly, the screw node in the BOMs of FIGS. 22 and 23 represent the same physical part. Accordingly, the BOM table 1217 only includes one record for the screw part.

Each record in the BOM table 1217 includes a BOM_id attribute which stores a key that is unique to the associated BOM node. Each record of the BOM table 1217 also includes a BOM_guid representing an alternate key of the record, a name attribute storing a name of the BOM node, a level_id attribute representing the logical level of the BOM node (that is, whether the BOM node represents an assembly, a subassembly, a part, etc.), a status_id attribute representing the status of the BOM node in the customer's business (that is, whether the part, subassembly, or assembly represented by the BOM node is in production, is discontinued, is contemplated for a future product, etc.), a creation_date attribute

identifying the date that the record was created, a `modification_date` attribute indicating the date that the BOM record was last modified, and a `version_number` attribute identifying the version of the BOM record.

The `level_id` attribute is actually a code. The values of the `level_id` codes are defined in a `BOM_level` table 1220. Similarly, the `status_id` attribute is a code whose values are defined in a `BOM_status` table 1221.

As evident by the examples of FIGS. 22 and 23, a BOM is a hierarchical and recursive structure, where each node in the structure may be a child of multiple parent nodes, and/or may itself have multiple child nodes. Information pertaining to this BOM hierarchical structure is not contained in the BOM table 1217. Instead, this information is contained in the `BOM_BOM_xref` table 1219.

The `BOM_BOM_xref` table 1219 includes a record for each parent/child relationship in the BOMs represented in the BOM table 1217. Each record of the `BOM_BOM_xref` table 1219 includes a `parent_BOM_id` attribute and a `BOM_id` attribute. The `parent_BOM_id` attribute stores the `BOM_id` of the parent BOM node, and the `BOM_id` attribute stores the `BOM_id` of the child BOM node. Each record of the `BOM_BOM_xref` table 1219 also includes a `sort_order` attribute which specifies the order in which the records in the `BOM_BOM_xref` table 1219 should be displayed or printed, when an appropriate command is received from the user. The value of the `sort_order` attribute is set to a default value, which is preferably 100. The user may change this value to any other number, either below or above 100. When printing or displaying the records of the `BOM_BOM_xref` table 1219, such records are printed or displayed in ascending or descending order (as indicated in the user command) based on the `sort_order` attribute.

The BOM table 1217 and the `BOM_BOM_xref` table 1219 shall be further described by reference to examples shown in FIGS. 24-25. FIG. 24 includes an example BOM table 1217X and FIG. 25 illustrates an example `BOM_BOM_xref` table 1219X. These tables in FIGS. 24 and 25 store data corresponding to the example BOMs 2202 and 2302 shown in FIGS. 22 and 23, respectively.

-141-

Accordingly, the example BOM table 1217X in FIG. 24 includes an entry for each BOM node in the bicycle BOM 2202 and the lawnmower BOM 2302. For simplicity purposes, only the BOM_id attribute and the name attribute are shown in the example BOM table 1217X. Accordingly, the BOM table 1217 includes an entry for the bicycle assembly, the bicycle frame subassembly, the screw part, the wheel subassembly, the rim part, the tire part, the bicycle handle bar subassembly, the metal rod part, the hand grip part, the lawnmower assembly, the lawnmower frame subassembly, the engine subassembly, and the lawnmower handle bar subassembly.

Information pertaining to the hierarchical structure of the bicycle BOM 2202 and the lawnmower 2302 is contained in an example BOM_BOM_xref table 1219X shown in FIG. 25. The BOM_BOM_xref table 1219X includes a record for each parent/child relationship in the BOM tables 2202 and 2302. Accordingly, record 2502 in the BOM_BOM_xref table 1219X indicates that BOM node 2 is a child of BOM node 1. In other words, record 2502 indicates that the bicycle frame node is a child of the bicycle node. Similarly, record 2504 indicates that the screw node is a child of the bicycle frame node, record 2506 indicates that the wheel node is a child of the bicycle node, record 2508 indicates that the rim node is a child of the wheel node, and record 2510 indicates that the tire node is a child of the wheel node. Further, record 2512 indicates that the bicycle handlebar node is a child of the bicycle node; record 2514 indicates that the metal rod node is a child of the bicycle handlebar node; record 2516 indicates that the handgrip node is a child of the bicycle handlebar node; record 2518 indicates that the lawnmower frame node is a child of the lawnmower node; record 2520 indicates that the screw node is a child of the lawnmower frame node; record 2522 indicates that the engine node is a child of the lawnmower node; record 2524 indicates that the lawnmower handle bar node is a child of the lawnmower node; and record 2526 indicates that the hand grip node is a child of the lawnmower handle bar node.

-142-

A BOM group represented in the BOM table 1217 may contain any number of documents. Preferably, the documents in a BOM group include patents that map to the BOM group. In other words, the documents in a BOM group include patents that map to the product represented by the BOM node corresponding to the BOM group in the BOM table 1217. As used herein, the term "product" refers to any assembly, subassembly, or part in a BOM.

A BOM_patent_xref table 1218 stores information that identifies the documents in a BOM group. The BOM_patent_xref table 1218 includes a record for each document in each BOM group. Each record of the BOM_patent_xref table 1218 includes a BOM_id attribute that identifies the BOM group, and a document_id attribute that identifies the document that is in that BOM group.

FIG. 26 illustrates an example BOM_patent_xref table 1218X that corresponds to the example in FIGS. 24 and 25. The example BOM_patent_xref table 1218X indicates that documents D1 and D2 are in the bicycle BOM group. Documents D1 and D2 are patents (see FIG. 17). Accordingly, the patents having document_ids D1 and D2 map to the BOM group having BOM_id 1 (that is, the bicycle assembly).

The preferred methodology for initially loading the BOM databases 626 is represented by a data flow diagram in FIG. 47 and a corresponding flowchart 9702 in FIG. 97. In step 9706, customer BOM data is received. Typically, a customer will have his own database that stores BOM data. In preparation for performing step 9706, the customer exports the BOM data from his database into a computer file. This computer file is represented as the customer BOM data 4704 in FIG. 47.

In step 9708, a BOM filter 4706 formats the customer BOM data 4704 so as to conform with the format of the BOM databases 626. The formatting done by the BOM filter 4706 may be minor or extensive, depending on whether, during the extract of the BOM data from the customer's databases, the customer was able to format the extracted data (represented as the customer BOM data 4704)

-143-

in a format that closely matched the format of the BOM databases 626. Database filters, such as the BOM filter 4706, for formatting data in preparation for upload to a database are well-known.

In some embodiments, the invention includes a filter that is specific to the database format of the BOM databases 626 (this is true for all filters discussed herein -- that is, the filters are specific to their respective databases). The structure and operation of this specialized filter will be apparent to persons skilled in the relevant art(s) based on the discussion herein, particularly the discussion of the database formats.

In step 9710, the formatted customer BOM data 4704 is loaded into a portion of the BOM databases 626. This portion is indicated as the first part 4708 in FIG. 47. It is not possible to load the remaining parts of the BOM databases 626, indicated as the second part 4710, using only the customer BOM data 4704, since the customer BOM data 4704 probably does not include the data attributes represented by the second part 4710. For example, the customer BOM data 4704 probably does not contain BOM logical level information or status information required by attributes in the BOM table 1217. Additionally, the customer BOM data 4704 will not include the patent mapping information (that is, an identification of which patents/documents map to which BOM groups) required by the BOM_patent_xref table 1218.

In step 9712, data on the customer's products is obtained. This data is obtained from the customer, such as the customer's production department or R&D department, and pertains to the status of the products (such as whether the products or parts are currently in production, whether they are contemplated for future production, whether they are discontinued, etc.).

In step 9714, patent related information is obtained. This patent related information may include, for example, a list of the customer's patents, and a list of competitor patents. Bibliographic information on these patents are loaded into the patent bibliographic databases 604 in the manner described above, to the

extent that the patent bibliographic databases 604 do not already include patent bibliographic information on these patents.

5 In step 9716, operators 4712 may run methodology reports which will be later used to aid in the analysis of the data gathered in prior steps. The methodology reports which may be run include a patent mapping report (described below). The patent mapping report may be incomplete due to the lack of information in the BOM databases 626. However, even an incomplete patent mapping report may be useful to aid in the analysis of the data.

10 In step 9718, operators 4712 analyze the customer BOM data 4704, the product data 4716, and the patent related information 4714 to determine, for each BOM group, the BOM level, the BOM status, and the documents (that is, patents) that map to the BOM groups. For example, the operators 4712 will analyze the BOM data and the product data to determine the status (that is, whether discontinued, in production, contemplated for future production, etc.) of the BOM node associated with each BOM group. Additionally, the operators 4712 will
15 analyze the customer BOM data 4704 and the patent related information to identify patents which map to the BOM group (that is, patents which map to the products represented by the BOM nodes corresponding to the BOM groups).

20 In step 9720, the information generated in step 9718 is loaded into the appropriate attribute fields of the BOM databases 626 (i.e., the second part 4710 of the BOM databases 626).

The steps of flowchart 9702 are periodically performed to update the BOM databases 626 with changes in BOMs, patents that map to BOMs, production changes, new products and features, etc.

25 In an alternate embodiment, the BOM databases 626 are not preloaded in the manner shown in FIG. 47 and discussed with reference to FIG. 97. Instead, referring to FIG. 48, BOM data is downloaded from a corporate BOM database 4808 via a corporate BOM database management system 4806 to the enterprise server 314 when needed. A translator 4804 translates the customer BOM data

-145-

from the format of the corporate BOM database 4808 to the required format of the BOM databases 626 in the enterprise server 314.

It is noted that this alternative embodiment could also be used with the other tables of the databases 316.

5 The approach of FIG. 48 is advantageous because the enterprise server 314 is assured of receiving and, accordingly, processing the most up to date BOM data. The disadvantage of the approach in FIG. 48 is that, typically, analysis of a customer corporation takes a number of weeks. During that period of analysis, the BOM data should be stable. In other words, during the period of analysis, the
10 analysis should be conducted on a snapshot of the customer's BOM data. The approach illustrated in FIG. 48 would not lend this stability, since the corporate BOM database 4808 will no doubt change during the course of the analysis period. Accordingly, the enterprise server 314 when downloading BOM data from the corporate BOM database 4808 will be working with different versions of the
15 BOM data at different times. In contrast, the approach of FIGS. 47 and 97 provide the desired stability of the BOM data during the analysis period.

 BOM groups (and, in fact, any predefined group) can also be represented using user-defined groups as long as the operator is willing to abstract and map the BOM group attributes to the attributes supported by the
20 user-defined groups. Such definition of BOM groups may be useful with BOMs that change frequently but core technologies change slowly, as in consumer products.

Corporate Entity Databases

 A corporate entity is considered to be a predefined group because any
25 number of documents can be associated with the corporate entity. These documents are preferably relevant and of interest to the corporate entity. For

-146-

example, the documents could be patents that are owned, licensed, or otherwise of interest to the corporate entity.

The corporate entity databases 630 include information on the customer corporation, the corporations associated or affiliated with the customer corporation, and/or corporations that are otherwise of interest to the customer corporation, such as the customer corporation's competitors or potential competitors.

The corporate entity databases 630 are shown in FIG. 12J. A corporate_entity table 1230 includes a record for each corporate entity of interest. In each record of the corporate_entity table 1230 includes a corp_entity_id attribute that stores a key that is unique to the corporate entity. Each record of the corporate_entity database 1230 also includes a corp_level_id attribute that specifies the level of the corporate entity in its respective corporate structure (it is noted that data pertaining to one or many corporate structures may be stored in the corporate_entity table 1230). Possible values of the corp_level_id attribute include parent corporation, subsidiary corporation, franchise, headquarters, etc. A name attribute stores the name of the corporate entity.

Information on the corporate structure is also stored in the corporate_entity table 1230. In particular, each entry of the corporate_entity table 1230 includes a parent_corp_entity_id attribute. This attribute stores the corp_entity_id of the immediate parent corporation of the corporate entity.

A sort_order attribute stores information that identifies the order in which the records of the corporate_entity database 1230 should be displayed or printed when requested by the user. The sort_order attribute is set to a default value of preferably 100. The user can change this value of the sort_order attribute. The records of the corporate_entity table 1230 are displayed and printed according to the sort_order attribute in the records. In particular, the records of the corporate_entity table 1230 are displayed or printed in ascending or descending order (specified by the user) of the sort_order attributes in the records.

-147-

The corp_level_id attribute is implemented as a code whose values are defined in a corporate_level database 1232.

5 A primary_corp database 1231 stores the corp_entity_id of the customer corporation. This primary_corp table 1231 is referenced when the enterprise server 314 needs to identify who the customer is during analysis operations. For example, if the user requests a listing of all patents owned by the customer corporation, the enterprise server 314 refers to the primary_corp table 1231 to identify who the customer is, and then refers to the corp_patents_xref table 1233 (described below) to identify all patents owned by the corporate customer.

10 A corporate entity group represented in the corporate_entity table 1230 can include any number of documents, such as any number of patents. Typically, the patents in a corporate entity group represent patents for which the corporate entity has some interest, such as patents owned by the corporate entity, patents assigned to the corporate entity, or patents that, for whatever reason, the corporate entity is analyzing. A corp_patent_xref table 1233 stores information that identifies the documents contained in each corporate entity group. The corp_patent_xref table 1233 includes a record for each document in each corporate entity group.

20 Each record of the corp_patent_xref table 1233 includes a corp_entity_id attribute that identifies the corporate entity group, and a document_id attribute that identifies the document that is in the corporate entity group.

25 Each record of the corp_patent_xref table 1233 also includes an ownership_id attribute and a relevance_id attribute. The ownership_attribute identifies the relationship of the document to the corporate entity. Possible values of the ownership_id attribute include owned, assigned, interested in, etc. The ownership_id attribute is implemented as a code whose values are defined by a patent_ownership table 1234, and are ultimately under the control of (i.e., defined by) the customer.

-148-

The relevance_id attribute identifies the relevance of the document to the corporate entity. Preferably, the possible values of the relevant_id attribute are core and non-core. A relevance of core indicates that the patent (identified by the document_id) maps to an assembly, subassembly or part that is currently in
5 production, or that is contemplated for future production, or for some other reason is important to the corporate entity. As discussed above, such assemblies, subassemblies and parts are represented in the BOM table 1217 as BOM groups. A non-core relevancy value indicates that the patent either does not map to any assembly, subassembly or part of interest or that the patent maps to an assembly,
10 subassembly, or part that is currently discontinued, or that, for whatever reason the corporate entity does not have great interest in the patent. The relevance_id attribute is implemented as a code whose values are specified in a relevance table 1235, and are ultimately under the control of (i.e., defined by) the customer.

The corporate entity databases 630 shall be further described with
15 reference to an example in FIGS. 33-36. These figures correspond to an example corporate structure 3202 illustrated in FIG. 32, wherein corp1 has two child corporations, corp2 and corp3, and corp3 has a child corporation, corp4. Corp5 is a competitor of Corp1.

FIG. 33 illustrates an example corporate_entity_table 1230X. For
20 simplicity purposes, only the corp_entity_id attribute, the corp_level_id attribute, and the parent_corp_entity_id attribute are shown in FIG. 33. The corporate_entity table 1230X has a record for each corporate entity of interest (that is, the corporate entities shown in FIG. 32). Each of these records in the corporate_entity table 1230X represents a corporate entity group that may
25 contain any number of documents. Record 3304 in the corporate_entity table 1230X indicates that corp2 is a subsidiary of corp1. Similarly, record 3306 indicates that corp3 is a subsidiary of corp1, and record 3308 indicates that corp4 is a subsidiary of corp3.

Note that corp5 in FIG. 32 is not a part of the corporate structure that includes corp1, corp2, corp3 and corp4. In the example of FIG. 32, corp5 is a competitor of corp1.

The corporate_entity table 1230X also includes a record 3310 for corp5. As indicated above, corp5 is a competitor of the customer corporation corp1. Therefore, the example of FIG. 33 indicates that the corporate_entity table 1230X includes a record for each corporate entity of interest including the competitors of the customer corporation corp1.

FIG. 36 illustrates an example corp_patent_xref table 1233X. This table indicates that the corp1 group includes documents D1, D2, and D3. These are patent documents (see FIG. 17). The ownership_id attribute and the relevance_id attribute are defined in the example patent_ownership table 1234X in FIG. 35 and in the example relevance table 1235X in FIG. 34, respectively. Pursuant to these tables, document D1 is indicated as being owned by corp1 and being a core patent. Similarly, documents D2 and D3 are indicated as being owned by corp1 and being core patents.

The methodology for initially loading the corporate entity databases 630 shall now be described with reference to the data flow diagrams in FIGS. 51 and 52 and the flowchart 10002 in FIG. 100. In step 10006, information on all corporate entities of interest is obtained. Such information includes compiling a list of the corporate entities of interest, determining their names, and determining their respective corporate structures. This information is represented in FIG. 51 as corporate entity data 5104.

In step 10008, a corporate entity filter 5106 formats the corporate entity data 5104 to conform to the format of the corporate_entity table 1230. Database filters for formatting data prior to upload to a database are well known.

In step 10010, the formatted corporate entity data 5104 is loaded into a part of the corporate_entity table 1230 (this portion is indicated as the first part 5108 in FIG. 51). It is possible that the corporate entity data 5104 may include

-150-

all the data necessary to load all of the fields of the corporate_entity database 1230. In this case, the formatted corporate entity data 5104 is also loaded into the second part 5110 of the corporate_entity database 1230. In other cases, however, the corporate entity data 5104 will not include all of the information necessary to load all of the fields of the corporate_entity table 1230, such as the parent_corp_entity_id attribute. In such cases, in step 10012, an operator 5112 analyzes the corporate entity data 5104 and determines values for the corp_level_id attribute and the parent_corp_entity_id attribute, when necessary. This information is loaded into the remaining portion of the corporate_entity table 1230 (this remaining portion is identified in FIG. 51 as the second part 5110).

In step 10014, information on patents of interest which are not already represented in the patent bibliographic databases 604 is obtained. This information is loaded into the patent bibliographic databases 604 in the manner discussed above to the extent that it is not already represented in the patent bibliographic databases 604.

Referring to FIG. 52, this patent information that is of interest to the corporate entities represented in the corporate_entity table 1230 is represented as patent information 5204.

In step 10016, operators 5208 map patents to corporate entities. In particular, these operators 5208 determine which patents are relevant to which corporate entities. This data is stored in the corporate_patent_xref table 1233. In other words, the operators 5208 determine which patents go into which corporate entity groups, and store this information in the corp_patent_xref table 233.

In step 10018, the operators 5208 and 5214 run methodology reports, as needed, such as patent mapping reports. These reports will aid the operators 5208 and 5214 in performing the analysis of 10020.

In step 10020, the operators 5208 and 5214 identify the ownership relationship for each patent/corporate entity pair specified in the corp_patent_xref

-151-

table 1233. In other words, for each record in the corp_patent_xref table 1233, the operators 5208 and 5214 identify the ownership relationship between the corporate entity and the patent. Additionally in step 10020, the operators 5208 and 5214 identify, for each record of the corp_patent_xref table 1233, the relevance of the patent to the corporate entity. Relevance can be core, non-core, etc.

In step 10022, this additional data generated in step 10020 is stored in the appropriate fields of the corp_patent_xref table 1233. These fields are presented by the second part 5212 in FIG. 52.

The steps of flowchart 10002 are periodically performed to update the corporate entity databases 630 to reflect changes in corporate structure, corporate acquisitions, corporate acquisition of patents, product line changes (that may change the relevance_id attribute), patent expirations (that would change the ownership_id attribute), etc.

Inventor, Employee, and Person Databases

An inventor is a predefined group because any number of patents can be associated with the inventor (i.e., patents where the person has been named an inventor). The inventorship databases 628 of the present invention preferably include information on the customer's past and present employee inventors, inventors of other companies, such as competitors, and inventors of any other patents of interest to the customer.

The inventor databases 628 are shown in FIG. 12J. These inventorship databases 628 interact with the person databases 632 and the employee 634. The person and employee databases 632 and 634 are shown in FIG. 12J and 12L.

In particular, the person table 1242 represents the person databases 632. The person table 1242 includes a record of each person of interest to the customer. These persons could be past and present employees of the customer,

inventors of patents of interest to the customer, employees (past and present) of competitors, etc. Each record of the person table 1242 includes a person_id attribute that represents a unique key of the person. Also included in each record of the person_id table 1242 are attributes that store the first name, middle name, and last name of the person, and the preferred suffix (Mr., Mrs., Ms., Dr., etc.) of the person.

An employee table 1243 is a table in the employee databases 634. The employee table 1243 includes a record of each employee of interest (whether or not the employee is an employee of the customer or a competitor, or is a past or present employee). Each record of the employee database 1243 includes a person_id attribute that identifies the person (from the person table 1242), and a corporate_entity_id attribute that identifies the corporate entity who is related to the person_id. There could be multiple entries in the employee table 1243 for the same person, if the person worked for many corporate entities represented in the databases 316.

Each record of the employee table 1243 further includes a job_function_id attribute that identifies the job function of the employee (for example, engineer, attorney, computer programmer, etc.), an employee_type_id that identifies the type of employee (such as part-time, full-time, etc.), an employee_status_id attribute that identifies the status of the employee (active, inactive, etc.), a start_date attribute and an end_date attribute that identify the period that the employee was employed by the corporate entity, and an employee_id attribute that indicates the employee's identification number within the corporate entity. The job_function_id attribute, the employee_type_id, and the employee_status_id attribute are preferably implemented as codes whose values are defined in an employee_job_func table 1245, an employee_type table 1244, and an employee_status table 1246, respectively.

Referring to FIG. 12J, a validated_inventor table 1236 represents the inventor databases 628. The validated_inventor table 1236 stores information that

identifies the patents to which a person is a named inventor. In particular, the validated_inventor table 1236 includes a record for each patent for which a person or employee is an inventor. Each record of the validated_inventor table 1236 includes a person_id attribute that identifies a person, and a document_id attribute that identifies a patent with which the person is a named inventor. Each record of the validated_inventor table 1236 also includes a corp_entity_id attribute that identifies a corporate entity. This corporate entity is the corporate entity that the person was associated with when the patent was filed or issued, or the corporate entity with whom the person was employed when the invention of the patent was invented. If the person was not employed by or associated with a company, this field is left empty.

The inventor databases 628, person databases 632, and employee databases 634 will now further be described with reference to the example shown in FIGS. 37-39. FIG. 37 illustrates an example person table 1242X. For simplicity purposes, only the person_id attribute is shown. As indicated by FIG. 37, there are five persons of interest to the customer.

FIG. 38 illustrates an example employee table 1243X. For simplicity purposes, only the person_id attribute, the corporate_entity_id attribute, and the employee_status_id attribute are shown. As indicated by the example in FIG. 38, person1 is an active employee of corp1. Similarly, person2 is an active employee of corp2, person4 is an inactive employee corp1 and person5 is an inactive employee of corp3.

FIG. 39 illustrates an example validated_inventor table 1236X. For simplicity purposes, only the person_id attribute and the document_id attribute are shown in FIG. 39. As evident from FIG. 39, person1 is a named inventor in the patents having document_ids D1-D3. Person3 is a named inventor in the patents having document_ids D1, D4, and D5. Person5 is a named inventor in the patent having document_id D6.

The preferred methodology for initially loading the person databases 632, the employee databases 634, and the inventor databases 628 shall now be described with reference to the data flow diagrams in FIGS. 49 and 50, and the flowcharts in FIGS. 98 and 99. FIGS. 49 and 98 corresponds to the extract and load procedures for the person databases 632 and the employee databases 634. In step 9806, information on the customer's past and present employees of interest are obtained, preferably from the customer's HR (Human Resources) or personnel department. This data is represented in FIG. 49 as the customer provided person and employee data 4904.

In step 9808, information on other people of interest (such as inventors of patents of interest or employees of competitors) is obtained. Such information is represented in FIG. 49 as other person and employee data 4906.

In step 9810, a person filter 4910 modifies the format of the customer provided person and employee data 4904 and the other person and employee data 4906 to match the format of the person databases 632 (that is, the person table 1242). Database filters for formatting data in preparation for database uploading are well-known.

In step 9812, the formatted customer provided person and employee data 4904 and the formatted other person and employee data 4906 are loaded into the person databases 632 (that is, the person table 1242).

In step 9814, an employee filter 4914 formats the customer provided person and employee data 4904 and the other person and employee data 4906 to match the format of the employee databases 634.

In step 9816, the formatted customer provided person and employee data 4904 and the formatted other person and employee data 4906 are loaded into the employee databases 634.

It is noted that in some cases, depending on the state of the data 4904, 4906, loading of the person databases 632 and the employee databases 634 may require the involvement of operators 4912, 4916.

-155-

FIGS. 50 and 99 pertain to the loading of the inventor databases 628 (that is, the validated_inventor database 1236). In step 9906, information is obtained on the patents of interest to the customer. Such patents of interest include the customer's own patents and patents of competitors.

5 In step 9908, bibliographic information on the patents obtained in step 9906 is loaded into the patent bibliographic databases 604 to the extent that such information is not already stored in the patent bibliographic databases 604.

10 In step 9910, operators 5008 analyze the patent information 5004 in conjunction with the person databases 632 and the employee databases 634 to map the patents obtained in step 9906 to the persons represented in the person databases 632 and in the employee databases 634. In other words, in step 9910, the operators 5008 identify, for each person represented in the person databases 632 and in the employee databases 634, the patents where that person is a named inventor.

15 In step 9912, the information generated in step 9910 is modified by a validated inventor filter 5006 to place the information in a form conforming with the format of the validated_inventor table 1236. Such formatted information is then loaded into the validated_inventor table 1236.

20 The steps of flowcharts 9802 in FIG. 98 and 9902 in FIG. 99 are periodically repeated to update the inventor databases 628, the person databases 632, and the employee databases 634.

Financial Databases

25 Example financial databases 638 are shown in FIG. 12M. These example financial databases include a BOM financial table 1247 and a patent licensing table 1248. The BOM financial table 1247 includes financial information pertaining to BOM groups. Each record of the BOM financial table includes a BOM_id attribute which identifies the BOM group, and also includes attributes that specify

-156-

the R&D dollars that the customer spent to develop the assembly, subassembly, or part corresponding to the BOM group. Each record of the BOM financial table 1247 also includes attributes that specify the gross revenue generated by that assembly, subassembly, or part corresponding to the BOM group, and the net revenue generated by that assembly, subassembly, or part.

The patent licensing table 1248 includes financial information pertaining to the licensing fees generated by patents of interest. Each record of the patent licensing table 1248 includes a document_id attribute that identifies a document (more specifically, that identifies a patent). Each record of the patent licensing table 1248 also includes attributes that identifies the licensee of the patent, the licensor of the patent, and the current licensing revenue generated by such licensing of the patent.

It is noted that the financial tables 1247 and 1248 shown in FIG. 12M are only examples of the financial databases 638 supported and envisioned by the present invention. The invention is adapted and intended to include other financial tables directed to other financial information of interest to the customer. Accordingly, tables 1247, 1248 of FIG. 12M are provided for purposes of illustration, and not limitation.

The financial databases 638 are generally loaded as described above with reference to FIGS. 45 and 95. The extract and load procedure represented in these FIGS. 45 and 95 may need to be modified to accommodate the particular financial information being processed. Such modifications will be apparent to persons skilled in the relevant art(s) based on the discussion herein.

Security Databases

The security databases 636 are shown in FIG. 12K. A user_table database 1238 includes a record for each user (that is, each person having rights to any of the data in the databases 316). Each record of the user_table database 1238

-157-

includes a user_id representing a key that uniquely identifies the user. Each record of the user_table 1238 also includes a user_name attribute that is an alternate key of the record, a user_full name attribute that stores the full name of the user, and a password attribute that stores the user's password.

5 The present invention supports a multi-level security methodology. A first level of this security methodology specifies that the creator of a data item has full access privileges with respect to the data item. For example, consider the group_table database 1227 in FIG. 12I. Each record of the group_table database 1227 includes an owner_user_id attribute which identifies the creator (i.e., the
10 owner) of the associated user-defined group. The owner_user_id attribute is implemented using codes whose values (user_id) are defined by the user_table database 1238 (FIG. 12K). Accordingly, the creator of a user-defined group as represented in a record of the group_table database 1227 has full access privileges with respect to that user-defined group.

15 A second level of the invention's security methodology defines that a user may have specific access privileges with respect to an object, called a secured object. Such secured objects include, for example, a note, a document and its notes, a document, and a database. Such secured objects also include a user-defined group.

20 This level of the security methodology as it relates to user-defined groups as secured objects is implemented using a group_user_xref table 1237. It is noted that other types of secured objects can be implemented in a similar manner.

Each record of the group_user_xref table 1237 includes a group_id attribute that identifies a user-defined group, and a user_id attributed that defines
25 a user. Each record of the group_user_xref table 1237 also includes a permission attribute which defines the access privileges that the user has with respect to the group. Possible access privileges include no access privileges, a read access privilege, a change access privilege (which allows the user to both read and modify the data item), and a delete access privilege (that allows the user to read, modify,

-158-

and delete the data item). Other access privileges could also be defined. Accordingly, a record in the group_user_xref table 1237 specifies that a person (represented by user_id) has specific privileges (indicated by the permission attribute) to a particular user-defined group (represented by the group_id).

5 A third level of the security methodology of the present invention specifies that a user group (such as a department) may have a particular access privilege with respect to a user-defined group. If a user is in that user group, then that user would also have that access privilege with respect to the user-defined group. This level of the security methodology is implemented using a grp_usr_grp_xref table
10 1240. Each record of this table 1240 includes a group_id attribute that identifies a user-defined group, and a user_group_id that identifies a user group. The user_group_id attribute is implemented as a code whose values are defined in a user_group table 1241. Each record of the grp_usr_grp_xref table 1240 also includes a permission attribute that defines the access privilege that the user group
15 (specified by the user_group_id attribute) has with respect to the user-defined group (specified by the group_id attribute). Preferably, a user group permission is overridden by a user's explicit permission, allowing the system administrator to give a user a higher or lower permission than the user group that the user is in.

20 A usr_grp_usr_xref table 1239 specifies the users who are in each user group. In particular, each record of this table 1239 includes a user_group_id attribute that identifies a user group, and a user_id attribute that identifies a user who is in that user group.

25 The security databases 636 will now be further described with reference to an example in FIGS. 27-31. FIG. 27 illustrates an example user_table 1238X. As evident by FIG. 27, three users have access privileges with respect to the data in databases 316, Mike, Jen, and Kathryn. Referring again to FIG. 19, a group_table database 1227X is shown wherein user 1 (Mike) is indicated as being the owner of groups GRP1-GRP4. User 2 (Jen) is indicated as being the owner of group GRP5. User 3 (Kathryn) is indicated as being the owner of group GRP6.

Accordingly, according to the security methodology of the present invention, user 1 (Mike) has complete access privileges (that is, the delete access privilege) with respect to groups GRP1-GRP4. User 2 (Jen) has delete access privileges with respect to group GRP5. User 3 (Kathryn) has delete access privileges with respect to group GRP6.

FIG. 28 illustrates an example group_user_xref table 1237X. This table indicates that user 2 has the change privilege with respect to user-defined group GRP1, and no privileges with respect to user-defined group GRP2.

FIG. 31 illustrates an example grp_user_grp_xref table 1240X. This table 1240X indicates that a user group with user_group_id UGRP2 has delete access privileges with respect to user-defined group GRP5. Similarly, a user group UGRP3 has no access privileges with respect to user-defined group GRP5, and user group UGRP3 has delete access privileges with respect to user-defined group GRP6. FIG. 29 illustrates an example user_group table 1241X.

FIG. 30 illustrates an example usr_grp_usr_xref table 1239X. This table 1239X indicates that user 2 is in user group UGRP2, user 1 is in user group UGRP3, and user 3 is in user group UGRP4. Accordingly, considering both usr_grp_usr_xref table 1239X and grp_usr_grp_xref table 1240X, user 2 (Jen) has delete privileges with respect to GRP5, user 1 (Mike) has delete privileges with respect to user-defined group GRP6, and user 3 (Kathryn) has no privileges with respect to user-defined group GRP5.

The various levels of the security methodology of the present invention do not work individually. Instead, these security methodology levels work together in a well defined and integrated manner. This integrated security methodology is implemented by the security module 402 in the enterprise server 314. The operation of the security module 402 when processing a request to access a data item (in particular, a user-defined group) is depicted in a flowchart 11002 in FIG. 110.

-160-

The security module 402 performs the steps of flowchart 11002 to determine whether a user who is requesting an operation involving a data item has sufficient security access privileges with respect to that data item. Preferably, all operations performed by the enterprise server 314 are security checked. In other
5 embodiments, only some operations performed by the enterprise server 314 are security checked. For example, operations involving reading patent documents are not security checked in some embodiments because patents are widely available public documents.

The user requesting the operation involving the data item is called the
10 requestor for reference purposes. For purposes of example, the data item that is involved in the operation is assumed to be a user-defined group having associated with it a record in the group_table database 1227 (FIG. 12I).

In step 11006, the security module 402 determines whether or not the
15 requestor is the owner of the requested user-defined group. The security module 402 performs step 11006 by comparing the requestor's user_id with the owner_user_id attribute in the record of the group_table database 1227 corresponding to the requested user-defined group. If it is determined that the requestor is the owner of the requested user-defined group, then the requestor is granted all access privileges (that is, the delete access privilege) with respect to the
20 requested user-defined group. Processing of flowchart 11002 is then complete. If it is determined that the requestor is not the owner of the requested user-defined group, then step 11008 is performed.

In step 11008, the security module 402 determines if the requestor has an
25 explicit access right with respect to the requested user-defined group. The security module 402 performs step 11008 by referring to the group_user_xref table 1237 (FIG. 12K). The security module 402 determines whether there is a record in the group_user_xref table 1237 having group_id equal to the key of the requested user-defined group, and also having the user_id attribute equal to the user_id of the requestor. If such a record exists in the group_user_xref table

-161-

1237, then the requestor is given access to the requested user-defined group in accordance with the access right specified by the permission attribute in the record of the group_user_xref table 1237. Processing of flowchart 11002 is then complete. If such a record does not exist in the group_user_xref table 1237, then
5 step 11010 is performed.

In step 11010, the security module 402 determines whether or not the requestor is in a user group that has an explicit access right with respect to the requested user-defined group. The security module 402 performs step 11010 by
10 accessing the usr_grp_usr_xref table 1239 to identify all groups that has the requestor as a member. The security module 402 then accesses the grp_usr_grp_xref table 1240 to determine whether any of the user groups in which the requestor is a member has an explicit access right with respect to the requested user-defined group. This is done by identifying a record in the grp_usr_grp_xref
15 table 1240 where the group_id attribute is equal to the identifier of the requested user-defined group, and the user_group_id attribute is equal to an identifier of a user group in which the requestor is a member. If such a record exists, then the requestor is given the explicit access right specified by the permission attribute in the record of the grp_usr_grp_xref table 1240.

Processing of flowchart 11002 is then complete. If such a record does not
20 exist in the grp_usr_grp_xref table 1240, then step 11012 is performed.

In step 11012, the security module 402 accesses a record in the grp_usr_grp_xref table 1240 corresponding to a user group called World, if such a record is available. All users are considered to be a member of the user group World. The requestor is then granted the access rights associated with the user
25 group World (again, if such a World record is available). Referring to the example in FIG. 31, the World user group has the access right of read. Accordingly, in step 11012, the security module 402 grants the requestor the read access right with respect to the requested user-defined group. Processing of flowchart 11002

-162-

is incomplete. Note that, if a World record is not available, the requester is given no access rights.

Some embodiments of the invention also include other levels of security. These levels can be in addition to, or in place of, one or more levels of security shown in FIG. 110. For example, the security module 402 before performing step 11006 may check to see if the requester has been explicitly revoked access to a particular portion of the database 316 that includes the data object. The requester may be revoked access to all of the data in the databases 316. As another example, the security module 402 before performing step 11006 may check to see if the requester is a super user, in which case the requester has complete access rights (i.e., the delete privilege) as to all data in the databases 316. These additional security levels can be implemented using tables in addition to the ones shown in FIG. 12K, or using modified variations of the tables shown in FIG. 12K. For example, additional fields can be added to the user_table 1238 that would indicate whether the user was a super user, or whether all access privileges for the user had been revoked as to a part of or all of the databases 316.

In an embodiment of the invention, the security module 402 preferably processes these security levels in the following order (other orderings of the security levels could also be used):

- (1) Has the requester been revoked access? If not, then:
- (2) Is the requester a super user? If not, then:
- (3) Is the requester the owner of the data item? (step 11006) If not, then:
- (4) Does the requester have explicit rights as to the data item? (step 11008) If not, then:
- (5) Is the requester in a user group that has explicit rights as to the data item? (step 11010) If not, then:

-163-

- (6) If a World group exists, then use the access rights associated with the World group. Otherwise, the requester is granted no rights.

It is noted that the security databases 636 shown in FIG. 12K and as described above apply only to the user-defined group databases 624. In other words, the security databases 636 only apply security measures as to the user-defined group databases 624. However, the invention is adapted and intended to extend the security of the security databases 636 and the security module 402 to all tables in the databases 316. The modifications necessary to the databases 316 and to the security module 402 to extend the security methodology of the present invention to all of the tables in the databases 316 will be apparent to persons skilled in the relevant art(s). For example, in order to extend the security methodology of the present invention to the BOM databases 626, it is only necessary to modify the BOM table 1217 to include an owner_user_id attribute. It would not be necessary to modify any of the security databases 636, as long as the group_id attribute could store identifiers of BOM groups.

In order to extent the security methodology to the notes databases 640, it would only be necessary to modify the notes databases 640 to include an owner_user_id. No modifications to the security databases 636 would be necessary, as long as the group_id attribute in the security databases 636 could be overloaded to store the identifiers of notes. Alternatively, additional security tables could be added to the security databases 636.

Only select persons have the ability to view and/or modify the data in the security databases 636. Such persons must have a high security level, and are typically system administrators. The client security module 702 provides a user interface at the clients 304, 306 that enable persons with sufficient security to view and/or modify the data in the security databases 636.

-164-

Some security information may be obtained from the underlying operating system, such as user IDs and passwords, and perhaps even access privileges. This information is then loaded into the security databases 636.

Enterprise Server and Client Functional Modules

5 The analysis modules 416 in the enterprise server 314 and the corresponding client analysis module 716 in the clients 304, 306 are described in detail below. FIG. 10 illustrates a block diagram of the analysis modules 416 in the enterprise server 314. The analysis modules 416 in FIG. 10 are not intended to be an exhaustive list of analysis modules that the invention is adapted or intended to support. The invention, instead, can support any analysis modules
10 capable of processing data in the databases 316 in a useful manner as judged by the customer. The scope of the invention is intended to include these modules.

For the most part, the analysis modules 416 perform functions in a patent-centric and a group-oriented manner. In other words, the analysis modules 416
15 perform functions on patent related information, with or without consideration of other information. Also, the analysis modules 416 perform functions on documents in groups.

Patent Mapping Module

20 The patent mapping module 1002 operates to identify all patents that are mapped to a product. As discussed above, each product (assembly, subassembly, part) of interest to the customer is represented by a node in a BOM. Each BOM node is represented by a BOM group. Each BOM group is represented by a record in the BOM table 1217 (FIG. 12G). Accordingly, the patent mapping module 1002 operates to identify all patents within a BOM group. Such patents
25 map to a product that is represented by the BOM group.

-165-

The operation of the patent mapping module 1002 is represented by a flowchart 8402 in FIG. 84. In step 8406, the patent mapping module 1002 receives from the operator (working at a client 304, 306, called the requesting client 304, 306) information that identifies one or more BOM groups (for reference purposes, these are called the operator-specified BOM groups). In an embodiment of the invention, this information identifies one or more BOM group names. The patent mapping module 1002 matches these BOM group names to BOM_ids by searching in the BOM table 1217 for matches of the operator-provided BOM names to the name attribute.

In step 8408, the patent mapping module 1002 identifies the patents contained in the BOM groups having the BOM_ids identified in step 8406. The patent mapping module 1002 performs step 8408 by searching in the BOM_patent_xref table 1218 for records having BOM_id attributes equal to the BOM_ids identified in step 8406.

In step 8410, the patent mapping module 1002 retrieves patent bibliographical information pertaining to the patents identified in step 8408. This is done by using the document_ids from the records of the BOM_patent_xref table 1218 identified in step 8408 as indexes into the patent bibliographic databases 604. This patent bibliographic information is transferred to the requesting client 304, 306.

In step 8412, the corresponding client analysis module 716 in the requesting client 304, 306 displays the received information. FIGS. 59 and 60 illustrate example display formats which the client analysis module 716 uses to display the patent data. In the display format 5902 of FIG. 59, the patent mapping data is sorted by BOM groups of logical level "subassembly". There are three BOM groups of logical level "subassembly" displayed in the example of FIG. 59, indicated by reference numbers 5904, 5906 and 5908. The BOM group of logical level "subassembly" having reference number 5904 is called "analog special effects controller subassembly". Each of its child BOM groups of logical level "part" are

-166-

displayed under this title 5904. For each of these child BOM groups of logical level "part", the part description and part number (corresponding to data from the name attribute in the BOM table 1217) are displayed. Also displayed is the patent number of a patent that maps to the part (obtained from the BOM_patent_xref table 1218), as well as the part status (obtained from the status_id attribute in the BOM table 1217), the ownership (obtained from the corp_patent_xref table 1233), and the relevancy status (obtained from the corp_patent_xref table 1233).

The display format 6002 in FIG. 60 is similar to the display format 5902 in FIG. 59. Patent mapping data in the display format 6002 is also organized according to the BOM logical level "subassembly". Child BOM groups of logical level "part" are displayed below the respective titles of the BOM groups of logical level "subassembly".

In some embodiments of the invention, the patent mapping module 1002 can also perform a document (as opposed to a patent) mapping function. This functionality is depicted in flowchart 8502 in FIG. 85. In step 8506, the patent mapping module 1002 receives from the operator (at the requesting client 304, 306) information that preferably identifies one or more group names. The patent mapping module 1002 matches these group names to group_ids by reference to the BOM table 1217.

In step 8508, the patent mapping module 1002 identifies the documents that are in the groups corresponding to the group_ids identified in step 8506. The patent mapping module 1002 performs step 8508 by reference to the BOM_patent_xref table 1218. These identified documents may include patents and non-patent documents.

In step 8510, the patent mapping module 1002 extracts bibliographic information pertaining to the documents identified in step 8508 from the appropriate document bibliographic databases 602. This information is transferred to the requesting client 304, 306.

-167-

In step 8512, the client analysis module 716 in the requesting client 304, 306 displays the information using display formats similar to that shown in FIGS. 59 and 60.

Patent Citation Module

5 The patent citation module 1004 operates to identify, for a particular patent (called the source patent), the patents which were cited during the prosecution of the selected patent (these patents are called citing patents). This is called a backwards citation report, because its focus is on looking backwards from the perspective of the source patent.

10 The patent citation module 1004 also performs a forward citation function. In performing the forward patent citation function, the patent citation module 1004 identifies, for a source patent, the patents in which the source patent was cited. Again, these patents are called citing patents for reference purposes.

15 A flowchart 8602 in FIG. 86 represents the operation of the patent citation module 1004 while performing the backward patent citation function.

 In step 8606, the patent citation module 1004 receives information from the operator (at a client 304, 306 called the requesting client 304, 306) that identifies one or more BOM groups. In an embodiment of the invention, the information provided by the operator identifies one or more group names corresponding to one or more BOM groups. The patent citation module 1004 identifies the BOM_ids for the BOM groups corresponding to the BOM group names by conducting a search through the BOM table 1217 using the BOM group names (provided by the operator) as an index.

20

 More generally, the patent citation module 1004 operates on any type of group, such as user-defined groups, inventor groups, corporate entity groups, etc. The operation of the patent citation module 1004 is discussed herein with reference to BOM groups for illustrative purposes only.

25

In step 8608, the patent citation module 1004 identifies the patents which are mapped to the BOM groups corresponding to the BOM_ids identified in step 8606. The patent citation module 1004 performs step 8608 by searching in the BOM_patent_xref table 1218 using the BOM_ids identified in step 8606 as indexes.

In step 8610, the patent citation module 1004 identifies the patents (called the citing patents) that are cited in the patents identified in step 8608 by reference to the PatentRef table 1208 (FIG. 12C). The patent citation module 1004 performs step 8610 by processing each patent identified in step 8608 in turn. In particular, the patent citation module 1004 selects one of the patents identified in step 8608 and uses the document_id of the selected patent as an index into the PatentRef table 1208. As a result of this operation, the patent citation module 1004 will identify a record in the PatentRef table 1208 for each patent that was cited in the selected patent. The patent citation module 1004 performs this operation for each of the patents identified in step 8608.

Also in step 8610, patent bibliographic information is retrieved from the patent bibliographic databases 604 for each of the citing patents.

Also in step 8610, the patent citation module 1004 forwards the results of the above operations to the requesting client 304, 306.

In step 8612, the client analysis module 716 in the requesting client 304, 306, displays the data received from the patent citation module 1004 to the operator. The client analysis module 716 displays this data in a format selected by the operator. FIGS. 61-65 depict example patent citation reporting formats. The operator can elect to view the patent citation information in any of these formats.

For example, if the operator chooses to view the patent citation information organized by BOM group, then the display format illustrated in FIG. 61 and/or FIG. 63 is used. In the display format 6102 of FIG. 61, the patent citation data is sorted by BOM group. A portion of the patent citation data for a

-169-

BOM group titled "digital video recording system" is shown in the example of FIG. 61. A table is displayed for each patent mapped to this BOM group. In the example of FIG. 61, such patents include U.S. Patent Nos. 5,327,235 and 5,359,428. These patents are called source patents, and they were identified in step 8608. Their patent numbers are displayed in the patent citation report in column 6106 of the tables. The titles of the source patents are displayed in column 6108 of the tables. The patents which are cited by the source patents are listed in column 6110 of the tables. These citing patents were identified in step 8610 of flowchart 8602. Bibliographic information on the citing patents is displayed in the remaining columns of the tables, such as the assignee information which is displayed in column 6112, title information which is displayed in 6114, and date information which is displayed in 6116. Other bibliographic information may be contained in the table and may be viewed by appropriate use of the navigation icons 6118.

The display format 6302 in FIG. 63 is similar to the display format 6102 in FIG. 61. The patent citation data in the format 6302 is organized by BOM group. Listed below each BOM group are the patents which are mapped to that BOM group. These patents are listed by their patent numbers in column 6304, and are called source patents. The patents which are cited in these source patents (called citing patents) are listed by their patent number in column 6306. Bibliographic information on these citing patents are displayed in the following columns of the tables.

Referring again to step 8612, if the operator elects to display the patent citation information sorted by selected patent characteristics, such as owned/license, then the display format 6402 in FIG. 64 is used. The display format 6402 includes a field 6404 which indicates whether owned or licensed patents are being displayed. If the operator elects to display only owned patents, then field 6404 would be modified to indicate this, and patent citation information on only patents owned by the customer (or any other corporate entity identified

-170-

by the operator) would be displayed in the patent citation report using format 6402.

Referring again to step 8612, if the operator elects to display the patent citation information in conjunction with ownership information, then display format 6202 in FIG. 62 is used. In display format 6202, the patent citation data is sorted by source patent number. For each source patent, the patents which cite the source patent (called citing patents) are displayed by patent number. Other patent bibliographic information on the citing patents is also displayed. In addition, ownership information (that is, whether or not the citing patent is owned, licensed, etc.) is also displayed in column 6204. This information was previously retrieved in step 8610 by the patent citation module 1004 by conducting a search of the corp_patent_xref table 1233 (FIG. 12J) using as indexes the document_ids of the citing patents (identified earlier in step 8610) and the corp_entity_id of a corporate entity specified by the operator. Accordingly, the ownership information displayed in column 6204 of display format 6202 can be relative to any operator-specified corporate entity, not just the customer corporate entity. The operator may have specified a corporate entity for purposes of data mining in the information received by the patent citation module 1004 in step 8606.

Referring again to step 8612 in FIG. 86, the patent citation report can be performed and displayed in a recursive fashion with an operator specified depth. The display of such a recursive patent citation function is shown in FIG. 65. In the example of FIG. 65, patent 1 is the source patent, and patents 2-4 are cited in patent 1 (that is, patents 2-4 are citing patents with respect to source patent 1). The patent citation report of FIG. 65 is multi-leveled. Accordingly, the display format of FIG. 65 indicates that patents 5-8 are cited in patent 2, patent 9 is cited in patent 3, and patents 10-12 are cited in patent 4. The example format 6502 shown in FIG. 65 is a two level patent citation report. The present invention can support patent citation reports of any operator specified level. Other display formats could be used by the invention to more effectively display the results of

-171-

multi-level patent citation reports, such as the well-known Xerox hyperbolic tree display format.

The patent citation module 1004 collects data for a multi-level patent citation report by repetitively performing the steps of flowchart 8602, which will
5 be apparent to persons skilled in the relevant arts. The level of the report desired by the operator could be contained in the information that the patent citation module 1004 receives in step 8606.

The forward patent citation function performed by the patent citation module 1004 is depicted in a flowchart 8702 shown in FIG. 87. In step 8706, the
10 patent citation module 1004 receives information from the operator at the requesting client 304, 306 that identifies one or more BOM groups. In one embodiment, the information received from the operator identifies one or more BOM names that correspond to the one or more BOM groups. Also in step 8706, the patent citation module 1004 identifies the BOM_ids of the BOM groups by
15 conducting a search in the BOM table 1217 using the BOM names as an index.

In step 8708, the patent citation module 1004 identifies the patents which are mapped to the BOM groups associated with the BOM_ids identified in step 8706. The patent citation module 1004 performs step 8708 by conducting a
20 search in the BOM_patent_xref table 1218 using the BOM_ids identified in step 8706 as indexes.

In step 8710, the patent citation module 1004 determines the patents that cite the patents identified in step 8708. The patent citation module 1004 performs
25 step 8710 by conducting a search through the PatentRef table 1208 using the document_ids corresponding to the citing patents identified in step 8708 as indexes. The patent citation module 1004 forwards the results of the above operation to the client analysis module 716 in the requesting client 304, 306.

In step 8712, the client analysis module 716 displays this information received from the patent citation module 1004 on the client display monitor 1122. The client analysis module 716 displays this information to the operator in any one

of a number of display formats, examples of which are shown in FIG. 61-65. These display formats are described above.

Patent Aging Module

5 The patent aging module 1006 operates to identify and analyze the remaining terms of patents that map to products, corporate entities, or inventors (that is, that are in BOM groups, corporate entity groups, or inventor groups that correspond to the products, corporate entities, or inventors). The patent aging module 1006 also operates to identify and analyze the remaining terms of patents that are in user-defined groups. The operation of the patent aging module 1006
10 is depicted in a flowchart 8802 shown in FIGS. 88A and 88B.

In step 8806, the patent aging module 1006 receives from the operator at the requesting client 304, 306, information that identifies one or more groups. These can be user-defined groups or predefined groups, or any combination of such groups. In an embodiment of the invention, the patent aging module 1006
15 receives information from the operator that identifies group names that correspond to groups. The operator can identify such groups while located at the client 304, 306 by selecting one or more groups from the group hierarchy user interface (FIG. 58). This is true in all cases herein where the user identifies groups to the enterprise server 314.

20 In some embodiments, the information may also have to identify the group types (i.e., BOM group, corporate entity group, inventor group, user-defined group, etc.). The patent aging module 1006 identifies the appropriate group IDs of these groups by reference to the appropriate group tables (that is, the BOM table 1217 for BOM groups, the corporate_entity table 1230 for corporate entity
25 groups, the validated_inventor table 1236 for inventor groups, the group_table database 1227 for user-defined groups, etc.).

-173-

In step 8808, the patent aging module 1006 identifies the patents which are contained in the groups corresponding to the group IDs identified in step 8806. The patent aging module 1006 performs step 8808 by searching through the appropriate patent xref tables, for example, the BOM_patent_xref table 1218, the group_document_xref table 1228, the corp_patent_xref table 1233, the validated_inventor table 1236, etc.

In step 8810, the patent aging module 1006 identifies the issue dates of the patents identified in step 8808 by reference to the patent table 1222 (FIG. 12H). The patent aging module 1006 does this by searching through the patent table 1222 using the document_ids identified in step 8808. In some embodiments, in step 8810, the patent aging module 1006 retrieves the expiration dates of the patents identified in step 8808 by reference to the calc_exp_date attribute in the patent table 1222.

In step 8812, to the extent that the expiration dates of the identified patents were not loaded in the patent table 1222, the patent aging module 1006 calculates the expiration dates of the identified patents in accordance with the existing patent laws of the appropriate jurisdiction. The patent aging module 1006 does this by using the patent issue dates. The patent aging module 1006 then forwards the results of the above operation to the client analysis module 716 in the requesting client 304, 306.

The client analysis module 716 then displays this information received from the patent aging module 1006 in an operator selected format. Such display of the patent aging information is described in the following steps.

In step 8814, the client analysis module 716 determines whether the operator requested to display the patent aging information in a standard report format. If the operator wishes to display the patent aging information in a standard report format, then the client analysis module 716 uses the display formats shown in FIG. 66 and/or FIG. 67 to display the patent aging information.

-174-

The display format 6602 of FIG. 66 is sorted by BOM group. The information for a BOM group titled "analog controller & imaging system patents" is shown in the example of FIG. 66. The display format 6602 displays the remaining patent term for the patents that map to the BOM group in a bar graph format.

The display format 6702 shown in FIG. 67 is similar to that of display format 6602 shown in FIG. 66.

In step 8818, the client analysis module 716 determines whether the operator wishes to display the patent information in a standard format in conjunction with ownership information. If the operator elects to display the patent aging information in a standard report in conjunction with ownership information, then example display format 6902 shown in FIG. 69 is used. It is noted that, in order to utilize the display format 6902, the patent aging module 1006 must have previously have obtained ownership information on the patents identified in step 8808 from the corp_patent_xref table 1233, and forwarded this information to the requesting client 304, 306. Such an operation is represented by step 8820.

The example display format 6902 shown in FIG. 69 is organized by ownership status relative to a corporate entity which the operator previously identified (this corporate entity could have been identified, for example, in the information received by the patent aging module 1006 in step 8806). Thus, the information shown in the display format 6902 can refer to patents owned and licensed by the customer or any other corporate entity, such as the customer's competitors.

In step 8824, the patent aging module 1006 determines whether or not the operator wishes to view the patent aging information sorted by group. If the operator wishes to view the patent aging information sorted by group, then in step 8826 the patent aging information is displayed using an example format 6802

-175-

shown in FIG. 68. The display format 6802 displays patent expiration information sorted by group.

In step 8828, the patent aging module 1006 determines whether the operator wishes to view the patent aging information sorted by organization level.

5 If the operator wishes to view the patent aging information sorted by organization level, then the patent aging information is displayed in step 8834 using an example format 7002 shown in FIG. 70. It is noted that format 7002 can be used only if

10 corporate_level_id information is obtained from the corporate_entity table 1230 for each patent identified in step 8808, and if this information is forwarded to the requesting client 304, 306. This operation is performed by the patent aging module 1006 in the enterprise server 314, and is represented by step 8830.

Additionally, the display format 7002 also requires that ownership information be obtained from the corp_patent_xref table 1233 by the patent aging module 1006, and that this information be forwarded to the requesting client 304, 306. Such

15 ownership information is obtained for each patent identified in step 8808 and for the corporate entity identified by the operator in the information preferably received by the patent aging module 1006 in step 8806. This operation is

20 In the display format 7002 of FIG. 70, the patent aging information is sorted by organization level (that is, by the information from the corp_level_id attribute retrieved from the corporate_entity table 1230). Additionally, ownership information relative to the specified corporate entity is displayed along with remaining patent term information.

25 In step 8836, the patent aging module 1006 determines whether the operator wishes to view the patent aging information organized by BOM level. If the operator wishes to view the patent aging information organized by BOM level, then the client analysis module 716 displays the patent aging information in a display format similar to that shown in FIG. 70, however instead of being sorted by organization level as currently shown in FIG. 70, the patent aging information

-176-

is sorted by BOM level. It is noted that, in order to display the patent aging information in this manner, the BOM level_id information must be retrieved from the BOM table 1217 for each of the patents identified in step 8808, and ownership information for these identified patents and for the operator-specified corporate entity must be obtained from the corp_patent_xref table 1233. Such operation is performed by the patent aging module 1006 in the enterprise server 314, and is represented by steps 8838 and 8840.

Patent Clustering and Bracketing Module

The patent clustering/bracketing module 1008 in the enterprise server 314 operates to identify and graphically represent potential relationships between a source patent and citing patents, where the citing patents are either cited in the source patent, or cite the source patent. Accordingly, the patent bracketing/clustering module performs a backwards operation, where it operates to identify potential relationships between a source patent and citing patents which were cited during the prosecution of the source patent. Additionally, the patent bracketing/clustering module 1008 also performs a forward function, where it attempts to identify potential relationships between a source patent and citing patents in which the source patent was cited.

The operation of the patent clustering/bracketing module 1008 when performing its backwards operation is represented by a flowchart 8902 in FIG. 89. In step 8906, the patent clustering/bracketing module 1008 receives information identifying a patent number from an operator at the requesting client 304,306. The patent clustering/bracketing module 1008 operates to identify the document_id of the patent having this patent number by using the patent number as an index in a search of the patent table 1222. This patent is called the source patent, for reference purposes.

In step 8908, the patent clustering/bracketing module 1008 identifies patents (called citing patents) that are cited by the source patent. The patent clustering/bracketing module 1008 performs step 8908 by using the document_id identified in step 8906 as an index during a search of the PatentRef table 1208.

5 In step 8910, the patent clustering/bracketing module 1008 identifies the ownership of the source patent and ownership of the citing patents by reference to the assignee table 1201 and/or the core_patent_xref table 1233. The patent clustering/bracketing module 1008 forwards the results of the above operations to the client analysis module 716 in the requesting client 304,306.

10 In step 8912, the client analysis module 716 displays the information received from the patent clustering/bracketing module 1008 in the client display unit 1122. The client analysis module 716 displays this information using any one of a number of formats, such as the example formats shown in FIGS. 71-73.

15 In an embodiment, the display format 7102 in FIG. 71 displays the patent bracketing/clustering information in a horizontal-oriented tree format (this is also called a connected graph format). An icon 7104 at the root of the tree represents the source patent. This icon is filled with a first pattern designated by reference number 7150 in key 7154. Icons 7106, 7108 and 7110 represent the citing patents.

20 In the display format 7102, links 7148 are used to represent the relationship between the source patent (represented by icon 7104) and the citing patents (represented by icons 7106, 7108 and 7110). In particular, the links 7148 indicate that the citing patterns were cited during the prosecution of the source patent 7104.

25 The icons 7106, 7108, 7110 corresponding to citing patents are filled with the same pattern 7150 of the icon 7104 corresponding to the source patent if the citing patents are owned by the same corporate entity as the source patent. In the example of FIG. 71, the icons 7106, 7108, 7110 corresponding to the citing patents are all filled with the first pattern 7150, which is the same that is used to

-178-

fill the icon 7104 corresponding to the source patent. Accordingly, the example in FIG. 71 indicates that all of the patents (both the source patent and the citing patents) are owned by the same corporate entity.

FIG. 71 illustrates a clustering scenario, because the corporate entity who owns the source patent and the citing patents has managed to cluster his patents in a given area of technology. Of course, this observation assumes that the source patent and the citing patents pertain to the same area of technology. However, this assumption is a reasonable one, given the relationship of the patents to one another (that is, given that the citing patents were cited during the prosecution of the source patent, and therefore can be presumed to be related to the source patent).

The example display format 7202 shown in FIG. 72 illustrates a scenario where the source patent 7204 is not owned by the same corporate entity as the citing patents 7206, 7208, 7210 and 7212. The scenario in FIG. 72 represents a potential bracketing situation, because the corporate entity who owns the source patent has not managed to obtain additional patents in the area of technology of the source patent. Instead, other corporate entities have obtained patents (represented by icons 7206, 7208, 7210 and 7212) in this area of technology. The source patent is potentially surrounded or bracketed by the citing patents.

The patent clustering/bracketing module 1008 is capable of performing a multi-level, recursive patent bracketing/clustering function. The patent clustering/bracketing module 1008 performs such a function by performing steps 8908 and 8910 for each of the citing patents. This operation by the patent clustering/bracketing module 1008 can continue to any operator specified level of recursion. An example display format generated by the client analysis module 716 on the client display unit 1122 is shown in FIG. 73. The source patent is represented by icon 7304, and has two citing patents represented by icons 7306 and 7312. The patent corresponding to 7306 has itself two citing patents, represented by icons 7308 and 7310. The patent represented by 7312 has a citing

-179-

patent represented by 7314, which itself has two citing patents represented by icons 7316 and 7318. The display in FIG. 73 represents a three-level patent bracketing/clustering scenario. This scenario in FIG. 73 also represents a hybrid clustering/bracketing situation. Note that the corporate entity who owns the root source patent (represented by icon 7304) also owns some of the citing patents in each level (that is, the patents represented by icons 7306, 7310, 7314 and 7316). However, there are some citing patents in the various levels which are not owned by this corporate entity.

The operation of the patent clustering/bracketing module 1008 when performing its forward operation is represented by a flowchart 9002 in FIG. 90. In step 9006, the patent clustering/bracketing module 1008 receives information from the operator that identifies a patent number. The patent clustering/bracketing module 1008 identifies the document_id that corresponds to a patent having this patent number by reference to the patent table 1222. This patent is called the source patent, for reference purposes.

In step 9008, the patent clustering/bracketing module 1008 identifies patents (called citing patents) in which the source patent is cited. The patent clustering/bracketing module 1008 performs step 9008 by reference to the PatentRef table 1208.

In step 9010, the patent clustering/bracketing module 1008 identifies the ownership of the source patent and the citing patents by reference to the assignee table 1201 and/or the core_patent_xref table 1233. The patent clustering/bracketing module 1008 sends the information obtained by the above operation to the client analysis module 716 at the requesting client 304, 306.

In step 9012, the client analysis module 716 displays this information on the client display unit 1122 using any of a number of display formats, such as the display formats shown in FIGS. 71-73. These display formats were described above.

-180-

It is noted that the operation of the patent clustering/bracketing module 1008, in performing the backwards operation (FIG. 89) and/or the forwards operation (FIG. 90), can be done on a group basis, instead of the individual patent basis as described above. In performing this functionality on a group basis, the patent clustering/bracketing module 1008 identifies all patents contained in the operator-specified group, and then identifies all citing patents of those patents in the operator-specified group. The client analysis module 716 displays the results of such group-oriented processing using any of a number of display formats, such as the formats shown in FIGS. 71-73, or in well-known formats suitable for displaying hierarchical data, such as the Xerox hyperbolic tree display format. It is noted that such display formats useful for displaying hierarchical structured data, such as the Xerox hyperbolic tree format, could also be used in the case of the individual patent operation represented by FIGS. 89 and 90.

Financial Module

The financial modules 1010 in the enterprise server 314 perform patent-centric and group-oriented processing of the data in the financial databases 638. Examples of the functions performed by the financial modules 1010 include determining the research and design (R&D) expenditures on a product or product line basis, determining the R&D expenditures per inventor or per employee on a product or product line basis, determining net licensing revenue on a product or product line basis, determining the number of patents issued on a product or product line basis, determining patent maintenance fees on a product or product line basis, determining market share on a product or product line basis, determining the tax rate on a product or product line basis, determining marketing costs on a product or product line basis, determining selling costs on a product or product line basis, determining the number of outstanding shares (P/E) on a product or product line basis, determining revenue on a product or product line

-181-

basis, determining cumulative product revenue on a product or product line basis, etc. The financial modules 1010 can also perform the above processing on a geographical region basis, or on a time basis.

Reference is made to the financial databases 1247 and 1248 shown in FIG. 12M for an example of the processing performed by the financial modules 1010. For example, the financial modules 1010 operate to determine the R&D dollars on a product basis (more particularly, the R&D dollars on a group basis, such as a BOM group basis). Given a operator-specified product (or an operator-specified group, such as a BOM group), the financial modules 1010 reference the BOM financial table 1247 using the associated BOM_id as a key to identify the record in the BOM financial table 1247 associated with the operator-specified BOM group. From that record, the financial modules 1010 can retrieve the values stored in the R&D dollars attribute. This value represents the R&D dollars spent in developing the product associated with the BOM group. The financial modules 1010 can also retrieve from that record the gross revenue and the net revenue generated by that product.

As another example, the financial modules 1010 can determine the licensing revenue associated with any patent owned by the company. This is done by referencing the patent licensing table 1248 using the document_id of the operator-specified patent as an index to identify the record(s) in the patent licensing table 1248 corresponding to that operator-specified patent. Once the record is found, the financial modules 1010 can retrieve information on the licensee, licensor, and license revenue related to the license of the patent from the licensor to the licensee. It is noted that a given patent may have multiple entries in the patent licensing table 1248, corresponding to each licensor/licensee combination.

The financial modules 1010 transfer the information generated from the above-described operation to the client analysis module 716 in the requesting

client 304, 306. The client analysis module 716 displays this information on the client display unit 1122, and enables the operator to manipulate the information.

Inventor Patent Count Module

5 The inventor patent count module 1012 in the enterprise server 314 operates to analyze patent inventor information to identify the top inventors for an operator-specified group. Top inventors are defined herein as being persons who most frequently are named as inventors on the patents in the group. The operation of the inventor patent count module is represented in the flowchart 9102 and FIG. 91.

10 In step 9106, the inventor patent count module 1012 receives from the operator at the requesting client 304, 306 information that identifies one or more groups. These can be any types of groups, such as BOM groups, corporate entity groups, operator-defined groups, etc. In an embodiment of the invention, the information received from the user identifies the group names of groups. These
15 groups are called the operator-specified groups for reference purposes. In some embodiments, this information may also identify the types of the groups (i.e., BOM group, corporate entity group, user-defined group, etc.).

The inventor patent count module 1012 identifies the group IDs of the operator-specified groups by using the group names as indexes in searches of the
20 appropriate group table (that is, the BOM table 1217, the group_table database 1227, and the corporate_entity table 1230).

In step 9108, the inventor patent count module 1012 identifies patents that are contained in the operator-specified groups. The inventor patent count module 1012 performs step 9108 by accessing the appropriate cross-reference tables using
25 the group keys (identified in step 9106) as indexes. Such cross-reference tables include the BOM_patent_xref table 1218, the group_document_xref table 1228, and the corp_patent_xref table 1233.

-183-

In step 9110, the inventor patent count module 1012 identifies the inventors of the patents identified in step 9108 by reference to the validated_inventor table 1236. The inventor patent count module 1012 accesses the validated_inventor table 1236 using the document_ids of the patents identified in step 9108 as indexes. The inventor patent count module 1012 processes the patent and inventorship information obtained as described above to determine, for each person, the number of patents (in the operator-specified group) in which the person is named as an inventor.

In step 9112, the inventor patent count module 1012 extracts additional pertinent information from the databases 316, such as each person's employment status from the employee table 1243, the relevance of the identified patents from the corp_patent_xref table 1233, and the corporate levels of the corporate entities that own the identified patents from the corporate_entity table 1230. The inventor patent count module 1012 forwards the information resulting from the above operation to the client analysis module 716 in the requesting client 304, 306.

In step 9114, the client analysis module 716 displays this data to the operator on the client display unit 1122. The client analysis module 716 displays the inventor patent count information in any of a number of display formats, such as display formats shown in FIGS. 74-77.

The sample displays in FIGS. 74 and 75 are similar, wherein the inventor patent count information is displayed in a bar graph format.

The display format 7602 in FIG. 76 displays the inventor patent count information in more of a spreadsheet or tabular format. This table lists, for a given group, the persons who are named as inventors of patents. Also listed are the number of patents in which each person is listed as inventor, as well as other information such as each person's employment status. Also listed in the display format 7602 of FIG. 76 is, for each person, the number of core patents and the number of non-core patents in which the person is named as an inventor.

The example display format 7702 in FIG. 77 is also a spreadsheet or tabular format sorted by organization level. The information contained in the format 7702 of FIG. 77 is similar to the information contained in the display format 7602 of FIG. 76.

5 *Inventor Employment Information Module*

The inventor employment information module 1014 in the enterprise server 314 operates to obtain information on persons who are named as inventors in patents that map to operator-specified groups. The operation of the inventor employment information module 1014 is represented by a flowchart 9202 in FIG. 92.

10 In step 9206, the inventor employment information module 1014 receives information from the operator at the requesting client 304, 306. This information identifies one or more groups, called operator-specified groups for reference purposes. Such groups may be BOM groups, user-defined groups, corporate entity groups, etc. In an embodiment of the invention, the information received from the operator identifies one or more group names that correspond to the operator-specified groups. The inventor employment information module 1014 identifies the group IDs of the operator-specified groups by reference to the appropriate group tables, such as the BOM table 1217, the corporate_entity table 1230, the group_table database 1227, etc. The inventor employment information module 1014 accesses these group tables using the group names as indexes.

15 In step 9208, the inventor employment information module 1014 identifies the patents which are contained in these operator-specified groups. These patents are called the identified patents for reference purposes.

20 The inventor employment information module 1014 performs step 9208 by accessing the appropriate cross-reference tables using the group IDs identified

-185-

in step 9206. These cross-reference tables include the BOM_patent_xref table 1218, the group_document_xref table 1228, the corp_patent_xref table 1233, etc.

5 In step 9210, the inventor employment information module 1014 identifies the inventors of the identified patents (identified in step 9208) by reference to the validated_inventor table 1236. The inventor employment information module 1014 accesses the validated_inventor table 1236 using the document IDs for the identified patents identified in step 9208.

10 In step 9212, the inventor employment information module 1014 extracts from the databases 316 information about the inventors. Such information is extracted by the inventor employment information module 1014 from the person table 1242 and the employee table 1243.

15 In step 9214, the inventor employment information module 1014 extracts other information for display in reports, such as the names of the corporate entities who own or license the identified patents. Such information is retrieved from at least the corporate_entity table 1230. The inventor employment information module 1014 transfers the information obtained from the above processing to the client analysis module 716 at the requesting client 304, 306.

20 In step 9216, the client analysis module 716 displays this information on the client display unit 1122 in a variety of operator-selected display formats, such as the example display formats shown in FIGS. 78-80.

The display format 7802 in FIG. 78 sorts the inventorship data by group. Information pertaining to the patents that map to a particular group is displayed in tabular format. The display format 7902 in FIG. 79 is similar to the display format 7802 in FIG. 78.

25 The display format 8002 in FIG. 80 organizes data according to the corporate entity that owns the identified patent. Information pertaining to the patents owned by each corporate entity is presented in a tabular format.

Exporting Patent Data Module

The exporting patent data module 1018 in the enterprise server 314 operates to copy data from the databases 316 to an operator-specified location. For example, the exporting patent data module 1018, pursuant to operator
5 command, can export a subset of the patent database 614 to an operator-specified location. The operator could then utilize this data locally, separate from the enterprise server 314. Operation of the exporting patent data module 1018 is represented by flowchart 9402 in FIG. 94.

In step 9406, the client analysis module 716 in the client 304, 306 enables
10 the operator to select which of the patents in the patent database 614 to export. The operator can select such patents by group, or individually, or both by group and individually. The client analysis module 716 transfers this operator request to the exporting patent data module 1018 in the enterprise server 314.

In step 9408, the exporting patent data module 1018 causes the document
15 storage and retrieval module 408 to retrieve information from the patent database 614 corresponding to the patents selected by the operator in step 9406. The enterprise server 314 transmits this information to the client 304, 306.

In step 9410, the client analysis module 716 copies the information
retrieved from the enterprise server 314 to a operator-specified location.

Importing Patent Data Module

As discussed above, the exporting patent data module 1018 operates to
copy a portion of the customer's patent repository (stored in the patent database
614) to a operator-specified location. The operator can then work locally with this
patent information. Eventually, the operator may want to update the databases 316
25 in the enterprise server 314 with any updates made to this subset of the patent
database 614.

The importing patent data module 1016 in the enterprise server 314 operates to perform this functionality. In particular, the importing patent data module 1016 imports to the databases 316 data that a operator has been working with locally, separate from the enterprise server 314. The operation of the
5 importing patent data module 1016 is represented by a flowchart 9302 in FIG. 93.

In step 9306, the client analysis module 716 in the client 304, 306 accesses patent information stored in an operator-specified location 9306 which is local to the client 304, 306. The client 304, 306 transfers this patent information to the enterprise server 314.

10 In step 9308, the importing patent data module 1016 receives this information from the client 304, 306 and loads the patent information into the databases 316, including the patent database 614 and the patent bibliographic databases 604.

Methodology Embodiments

15 The present invention can be used to perform many patent-centric and/or group-oriented tasks. These tasks would be of interest to business customers. These tasks include, but are not limited to:

- determining patent ownership;
- mapping patents to products;
- 20 • mapping patents to corporate entities;
- mapping patents to employees;
- mapping employees to inventors;
- identifying patents that are licensed;
- identifying remaining patent terms on licensed patents;

-188-

- analyzing unused licensed patents to determine whether the technology covered by these patents would be helpful in developing new products or features;
- determining if licensed patents are being used by the company;
- 5 • determining the extent and the length of patent protection on core products;
- determining whether additional patent protection should be sought on core products;
- 10 • determining whether the company has adequately protected an area of technology;
- determining whether competitor companies have been able to protect portions of an area of technology;
- examining a competitor's patents on a product line basis;
- examining a competitor's patents via patent term analysis;
- 15 • examining a competitor's inventors;
- identifying potential infringement of the company's patents;
- determining whether key inventors are still employed by the company;
- determining whether inventors currently employed by competitors
20 should be hired;
- determining the extent to which a contemplated future product is currently protected by the company's patents;
- determining whether and to what extent patent applications should be filed on a contemplated future product;
- 25 • determining whether features of a contemplated future product are covered by competitors' patents;
- determining whether a present or future product should be modified in order to design around a competitor's patent;

-189-

- determining whether features in current or future products have been adequately protected by patents;
- determining whether the company should retain, sell, or expand a business area;
- 5 • determining whether non-core patents should be sold or licensed;
- determining whether the maintenance fees on non-core patents should be paid;
- determining whether the customer's human resources are being most effectively used;
- 10 • determining whether licensed patents cover the company's products in order to decide whether to maintain or cancel the licenses;
- determining whether the company should file continuation applications to offset the expiration of current patents;
- 15 • determining the R&D expenditures on a product basis;
- determining the R&D expenditures per inventor/employee on a product basis;
- determining the net licensing revenue on a product basis;
- determining the patent budget on a product basis;
- 20 • determining the number of patents issued on a product or product line basis;
- determining the patent maintenance fees on a product or product line basis;
- determining the customer's market share on a product or product line basis;
- 25 • determining the customer's tax rate on a product or product line basis;
- determining the marketing cost on a product or product line basis;
- determining the selling cost on a product or product line basis;

-190-

- determining the allocated cost on a product or product line basis;
- determining the number of outstanding shares (P/E) on a product or product line basis;
- 5 • determining the product revenue on a product or product line basis;
- determining the cumulative product revenue on a product or product line basis;
- determining the number of core patents on a product or product line basis;
- 10 • determining the number of non-core patents on a product or product line basis;
- determining the ratio of core to non-core patents on a product or product line basis; and
- determining the number of licensed patents on a product or product line basis.
- 15

The tasks listed above, and other tasks that would be useful to the business customer, can be performed by invoking the analysis modules 416, either individually or in combination. For example, some of the tasks mentioned above can be performed by invoking the patent mapping module 1002 alone. Others of

20 the tasks listed above require, for example, the invocation of the patent mapping module 1002 in combination with one or more of the other analysis modules 416, such as the patent citation module 1004 or the patent aging module 1006.

The performance of some of the tasks mentioned above involve human intervention. Specifically, with some of the tasks, one or more of the analysis

25 modules 416 are invoked. These analysis modules 416 automatically process data in the databases 316, and automatically prepare reports (generally called methodology reports) based on such processing. One or more human operators then analyze the reports to complete the tasks. The invention supports and assists

-191-

the human operators in performing their functions by automatically analyzing data and running reports.

In some embodiments, the invention includes expert system(s) 1020 to automatically analyze and process the data in the reports that were automatically generated by the analysis modules 416. It is noted that the expert system 1020 is shown as being a part of the analysis modules 416 in FIG. 10, but may in practice be a separate component of the enterprise server 314, or may form part of the client 304, 306.

In these embodiments, the expert system 1020 either aids the human operators in performing the above-listed tasks, or completely replaces the human operators in analyzing the data that was automatically generated by the analysis modules 416. Such expert systems 1020 could be trained (using well known procedures) with the knowledge and experience of the human operators. The structure and operation of expert and rule based systems, and the training of such systems, are well known and are described in a number of publicly available documents, including (but not limited to) Charniak et al., *Introduction to Artificial Intelligence*, Addison Wesley, 1986, and Nils J. Nilsson, *Principles of Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA, 1980, which are incorporated herein by reference in their entireties. The manner in which to construct and train expert systems to automatically perform a part of or all of the processing described herein as being performed by human operators will be apparent to persons skilled in the relevant art(s).

Various patent-centric and group-oriented methodologies of the present invention are described below. These methodologies represent tasks that can be performed (either with or without human involvement) through use of the invention. These methodologies essentially involve two steps. In the first step, one or more analysis modules 416 are invoked for the purpose of automatically accessing and processing data contained in the databases 316, and for the purpose of automatically reporting on such processing (such reports are generally called

-192-

methodology reports). In the second step, the reports generated by the analysis modules 416 are analyzed, and business-related actions, plans, and/or strategies are taken or developed based on the analysis. This second step may be performed by human operators, or may be automatically performed by expert systems (as described above), or may be performed by a combination of human operators and expert systems. In the following description of the methodologies, the second step is described as being performed by human operators. It should be understood, however, that the second step could alternatively be performed by expert systems, or by expert systems in combination with human operators.

The methodologies described below represent a sampling of the tasks that can be performed through use of the invention. The invention is adapted and intended to perform other patent-centric and group-oriented tasks, particularly tasks that can be implemented through mining and/or processing of the data in the databases 316. Accordingly, the following methodologies are described herein for purposes of illustration, and not limitation.

Patent Mapping and Mining

A corporate entity can utilize the present invention in order to perform a patent mapping and mining process. This process is an internally focused analysis where the impact of the corporate entity's patents (both owned and licensed) on the corporate entity is analyzed. The business impact of this process is to provide high financial leverage within the corporate entity. The patent mapping and mining process provides a tangible way of defining core (active) and non-core (non-active) patents and how they relate to the corporate entity's business. Short term financial results of the patent mapping and mining process can be found in licensing opportunities and non-renewal of inactive patents and licensed patents.

It is noted that the customer may, based on its own practices or needs, define the meaning of core and non-core. For example, core could correspond to

patents being used by the customer, and non-core could correspond to patents not being used by the customer. Accordingly, the definition of core and non-core is implementation dependent.

5 A flowchart 10202 in FIG. 102 illustrates an example patent mapping and mining process. In FIG. 10206, an operator (that is, an employee or agent of the customer corporate entity) at a client 304, 306 issues a command that causes the patent mapping module 1002 to run a patent mapping report with respect to a product specified by the operator. The patent mapping module 1002 runs a patent mapping report on the BOM group corresponding to the operator-specified product in the manner discussed above. Example patent mapping reports are shown in FIGS. 59 and 60.

10 In step 10208, the operator identifies from the patent mapping reports the core and non-core patents that map to the BOM group corresponding to the operator-specified product. For example, in the example display of FIG. 59, if the operator specified product is the analog special effects controller subassembly 5904, then the operator would be able to determine from the patent mapping report that U.S. Patent Nos. 4,719,523; 4,686,590; 4,716,476 and 5,029,013 were non-core patents. In contrast, if the operator specified group was the digital color correction subassembly 5906, then the operator could identify from the patent mapping report that U.S. Patent Nos. that 5,381,175; 5,452,018; and 5,353,059 were core patents.

15 In step 10210, the operator considers various business options with respect to the non-core patents. In particular, the operator considers whether the non-core patents should be licensed or sold. This determination may be based on whether the corporate entity is contemplating the introduction of future products having features that would be covered by the non-core patents.

25 The operator may also consider whether or not the non-core patents should be held in reserve as a shield against a possible allegation of infringement by a competitor. This issue may depend on whether the non-core patents cover

technology which is used by competitors. The operator can run patent citation reports (performed by the patent citation module 1004) and/or patent clustering/bracketing reports (performed by the patent clustering/bracketing module 1008) to determine whether it appears that its competitors are active in the areas of technology covered by the non-core patents.

The operator could also consider whether patent maintenance fees should be paid on the non-core patents. This issue, like those discussed above, may depend on whether the corporate entity intends to use the technology covered by the non-core patents in future products, or whether or not the corporate entity's competitors are currently using the technology covered by the non-core patents. A decision to discontinue payment of maintenance fees on the non-core patents would result in an immediate monetary savings for the corporate entity.

Also, the operator could consider whether or not the corporate entity should continue to license the non-core patents. The issues involved with this determination are similar to those discussed above. If the corporate entity decides to discontinue the license of the non-core patents, then the corporate entity could realize an immediate cost savings.

Further, the operator could consider whether the technology covered by the non-core patents could be used in the development of new features or products for the corporate entity.

In the above, step 10210 is described as being performed by a single operator. In practice, the performance of step 10210 may require the involvement of many people with knowledge in many different fields and having different abilities (or expert systems with this knowledge and abilities).

The patent mapping and mining process as represented in flowchart 10202 is an iterative process, as represented by control arrow 10212. In other words, the patent mapping and mining process represented in flowchart 10202 could be performed repeatedly over time for the same product, or could be performed a number of times, each time for a different product.

Situation Assessment

A corporate entity could use the present invention to perform a situation assessment process. A situation assessment process focuses on the corporate entity's current business situation. The situation assessment process determines the strategic "playing field" of the corporate entity's business as it relates to its patents and products. By performing a situation assessment process or analysis, the corporate entity can begin to analyze its current market in the competitive situation in relation to its products and patents. The situation assessment analysis begins by running a number of methodology reports in order to better understand the corporate entity's competitive situation. These methodology reports provide a clear understanding of a number of patent related business metrics, including:

- potential patent infringement;
- the length of protection on core product patents based on patent aging reports;
- how well the corporate entity's patents are clustered to protect current and future products and key features;
- determination of whether competitors are bracketing the company's current and future products through their patents;
- determination of whether key technologists who have developed the current core products are still in the company or have gone to competitors;
- a determination of whether key inventors and other R&D personnel are working on appropriate products based on the corporate entity's strategic direction; and
- a determination of whether any R&D resources should be reallocated for best use of human resources.

Based on the above patent related business metrics, the corporate entity can determine what strategic steps need to be implemented in relation to: merger and acquisitions, human resources (HR) retention programs, patent licensing programs, patent acquisition programs, product or division sale, appropriate tax strategies, etc.

A flowchart 10302 in FIG. 103 illustrates an example situation assessment process.

In step 10306, a competitive analysis process is performed to identify and analyze the positions of the corporate entity's competitors. In performing the competitive analysis process, the corporate entity examines the competitor's products per product line, examines competitor's patents via aging reports, profiles competitor's inventors, and determines whether competitors are potentially infringing the corporate entity's patents. Details of the competitive analysis process is described below.

In step 10308, patent clustering and bracketing are evaluated. In performing this evaluation, the corporate entity runs clustering and bracketing reports for future product protection, determines if product features have been adequately protected, determines if the corporate entity's present and future core technologies have been adequately protected by patents, etc. Details of the clustering and bracketing process are described below.

In step 10310, an inventor analysis is performed. In performing this inventor analysis, the corporate entity analyzes inventorship with respect to patents covering its core technology, analyzes competitor's inventors, determines whether it should consider hiring the inventors of its competitors, determines whether its own employee retention program should be modified in any way, etc. Details of the inventor analysis are described below.

In step 10312, the corporate entity performs a financial analysis process using the present invention. In performing this financial analysis, the corporate

-197-

entity evaluates a number of patent-centric and group-oriented business metrics, such as the R&D expenditures per product, the R&D expenditures per inventor for a product line, the net licensing revenue per product line, the patent budget for a product line, the patent maintenance fees associated with a product line, the market share of a product line, the tax rate of a product line, marketing costs of a product line, selling costs of a product line, allocated costs of a product line, product revenue over product costs, etc. Details of this financial analysis process are described below.

The situation assessment process is an iterative one, as represented by control arrow 10314 in the flowchart 10302. The steps of flowchart 10302 can be performed repeatedly, and in any order. Thus, the steps in FIG. 103 are provided for purposes of illustration, not limitation.

Competitive Analysis

FIG. 104 is a flowchart 10402 representing the operation of the competitive analysis process.

In step 10406, an operator associated with the corporate entity (such as an employee or consultant of the corporate entity) located at a client 304, 306 issues a command to the patent citation module 1004 to run a patent citation report on the patents in a group specified by the operator. The operator may request that either a backwards and/or a forwards patent citation report be formed. The manner in which the patent citation module 1004 performs the patent citation process is described above. Examples of patent citation reports are shown in FIG. 61-65.

In step 10408, the operator, by reference to the patent citation report, identifies the companies who are the current assignees of the citing patents. These companies may be potentially working in the same technological areas as the source patents since their patents either cited the source patents (in the case of

-198-

forwards patent citation reports), or they were cited in the source patents (in the case of backwards patent citation reports). Accordingly, the assignees of the citing patents may potentially be competitors of the customer corporate entity (assuming, of course, that the assignees of the citing patents are not the same as the customer corporate entity). Accordingly, in step 10408, the operator, through analysis of the patent citation report, is able to identify potential competitors of the customer client entity with respect to the technological area of the operator-specified group. It is noted that these potential competitors may be infringing the corporate entity's patents (that is, the source patents). These potential competitors may also be interested in licensing the corporate entity's patents (the source patents). Further, the customer corporate entity may be interested in licensing or purchasing the patents listed as citing patents in the patent citation report.

In step 10410, the operator further analyzes the patent citation report to determine the potential strength of the competitors' positions in the technological areas relative to the source patents. For example, if a large number of the citing patents are assigned to a given competitor company, then that may indicate that the competitor company has a strong patent position with respect to the technological area of the associated source patent. An operator may conclude from this that the customer corporate entity should consider reducing its product line or even discontinuing the product line given the potential strengths of the competitor's patents in the area. Other options would be to sell the product line to the potential competitor, or modify the product line to avoid any potential patent infringement by the customer client entity of the competitor's patents.

In step 10412, the operator reviews the patent citation report to identify a competitor's potential future product direction. The operator performs step 10412 by noting the number of patents that a potential competitor has in an area, and also noting the filing and issue dates of those patents. This information may indicate the extent to which a potential competitor is interested in a technological

area. For example, if a potential competitor has a large number of patents in an area, and those patents were filed or issued relatively recently, then that may be an indication that the potential competitor is preparing to move into or expand its presence in the area.

5 In step 10414, the operator issues a command that causes the patent aging module 1006 to run an aging report. The patent aging reports are run with respect to the citing patents (that is, the patents which are owned by potential competitors). The manner in which the patent aging module 1006 performs the patent aging function is described above. Examples of patent aging reports are
10 shown in FIGS. 66-70. The operator analyzes the patent aging report in view of issues including whether or not to license competitor patents, purchase competitor patents, design around competitor patents, ignore competitor patents, etc. These decisions are impacted by the remaining terms of the competitor patents. For example, if a potentially relevant competitor patent has a long patent term
15 remaining, then the customer corporate entity may be more inclined to license or purchase the patent. In contrast, if the patent has a short term remaining, then the customer corporate entity may be more inclined to ignore the patent or design around the patent.

20 In step 10416, the operator issues commands that cause the inventor patent count module 1012 and the inventor employment information module 1014 to run inventorship reports on the patents owned by the potential competitors. The manner in which the inventor patent count module 1012 and the inventor employment information module 1014 operate to run these inventorship reports are described above. Examples of inventor patent count reports are shown in
25 FIGS. 74-77. Examples of inventor employment information reports are shown in FIGS. 78-80.

 In step 10418, the operator analyzes inventor patent count reports to identify the top inventors of patents covering key technology. The operator analyzes the inventor employment information reports to determine if these key

-200-

inventors are still with the competitors. If they are not, then the operator can consider whether the corporate entity should consider hiring the inventors. Also from the inventor employment information reports, the operator can determine if the top inventors were former employees of the customer corporate entity. If they were, then the operator can set in motion processes to insure that these inventors are complying with their former employment agreements with the customer corporate entity. Further from the inventor employment information reports, the operator can determine if the top inventors are still employed by the competitors. If they are, the operator can consider whether the corporate entity should consider hiring these top inventors.

The competitive analysis process is an iterative one, as indicated by control arrow 10420. The steps of flowchart 10402 can be performed repeatedly, and in any order. Therefore, the steps of flowchart 10402 are provided in FIG. 104 for purposes of illustration, and not limitation.

The above operation is described as being performed by a single operator. In practice, the above operation may require the involvement of many people with knowledge in many different fields and having different abilities (or expert systems with this knowledge and abilities).

Clustering and/or Bracketing

The clustering and/or bracketing process is illustrated in a flowchart 10502 in FIG. 105. In step 10506, an operator associated with the customer corporate entity at a client 304, 306 issues commands to run patent clustering/bracketing reports for a particular product, called the operator specified product for reference purposes. This product corresponds to a BOM group. The patent clustering/bracketing module 1008 in the enterprise server 314 runs these patent clustering/bracketing reports for the operator specified BOM group. The

-201-

operation of the patent clustering/bracketing module 1008 is described above. Examples of patent clustering/bracketing reports are shown in FIGS. 89 and 90.

In step 10508, the operator studies the patent clustering/bracketing reports to determine whether the customer client entity's product has been adequately protected by the customer corporate entity's patents. In other words, the operator reviews the clustering/bracketing report to determine whether or not clustering exists. An example of clustering is shown in FIG. 71.

In step 10510, the operator reviews the patent clustering/bracketing reports to determine the extent to which others have patented the technology related to the patents in the BOM group associated with the operator-specified product. In other words, the operator in step 10510 analyzes the patent clustering/bracketing reports to determine whether full or partial bracketing exists. An example of full bracketing is shown in FIG. 72. An example of partial bracketing is shown in FIG. 73.

Based on the above analysis, the operator evaluates whether these patents owned by others should be licensed, purchased, ignored, designed around, etc. To assist in performing this evaluation, the operator may run patent aging reports and/or inventorship reports on these patents owned by others. These reports will aid the operator in conducting the evaluation.

The clustering and/or bracketing process is an iterative one, as represented by control flow arrow 10512. The steps of the clustering and/or bracketing flowchart 10502 can be performed repeatedly for the same or different BOM groups, and can be performed in any order. Accordingly, the steps of the clustering and/or bracketing flowchart 10502 are presented in FIG. 105 for purposes of illustration, and not limitation.

The above operation is described as being performed by a single operator. In practice, the above operation may require the involvement of many people with knowledge in many different fields and having different abilities (or expert systems with this knowledge and abilities).

-202-

Inventor Analysis

The inventor analysis process is illustrated in the flowchart 10602 shown in FIG. 106. In step 10606, an operator associated with the customer client entity located at a client 304, 306 issues commands to run inventorship reports for a product, called an operator specified product for reference purposes. The product corresponds to a BOM group. For references purposes, the BOM group is called the operator specified BOM group (because it corresponds to the operator specified product). The inventor patent count module 1012 and the inventor employment information module 1014 in the enterprise server 314 perform these inventorship report functions. In particular, the inventor patent count module 1012 performs the inventor patent count function, and the inventor employment information module 1014 performs the inventor employment information function. The operation of the inventor patent count module 1012 and the inventor employment information 1014 is described above. Examples of inventor patent count reports are shown in FIG. 74-77. Examples of inventor employment information reports are shown in FIGS. 78-80.

In step 10608, the operator identifies key or top inventors of the customer corporate entity's patents based on the inventorship reports.

In step 10610, the operator refers to the inventor employment information reports to identify whether these key inventors are still employed by the customer corporate entity. If they are still employed by the customer corporate entity, then the operator can initiate steps to ensure that the customer corporate entity retains them. If they are not still employed by the customer corporate entity, then the operator can initiate steps to ensure that these former employees are in compliance with their employment agreement with the customer corporate entity. Also, the operator can consider whether the customer corporate entity should consider rehiring them.

-203-

In step 10612, if the key inventors are still employed by the customer corporate entity, then the operator refers to the inventor employment information reports to determine whether the talents of these inventors are being most effectively utilized by the corporate entity. For example, if the inventors are no longer assigned to the department that developed the operator specified product, then the corporate entity may want to transfer them back to that department so that they can further enhance the product.

The inventor analysis process is an iterative one, as represented by control flow arrow 10614. The steps of the inventor analysis flowchart 10602 can be performed repeatedly or in any order for any given BOM group. Accordingly, the steps in the inventor analysis flowchart 10602 are provided for purposes of illustration, not limitation.

The above operation is described as being performed by a single operator. In practice, the above operation may require the involvement of many people with knowledge in many different fields and having different abilities (or expert systems with this knowledge and abilities).

Financial Analysis

The financial analysis process is represented by a flowchart 10702 in FIG. 107.

In step 10706, an operator associated with the customer corporate entity at a client 304, 306 issues instructions to the financial modules 1010 in the enterprise server 314. These instructions command the financial modules 1010 to calculate the research and development dollars per patent in a BOM group. The BOM group is specified by the operator.

The financial modules 1010 perform this function by first determining the number of customer patents that map to the BOM group (by reference to the BOM_patent_xref table 1218). Next, the financial modules 1010 determine the

-204-

R&D dollars for a project (by reference to the BOM financial table 1247). Then, the financial modules 1010 divide the R&D dollars by the number of patents.

In step 10710, the financial modules 1010 calculate the licensing dollars resulting from a project by identifying the customer patents that map to the product (by reference to the BOM_patent_xref table 1218), and determining the licensing revenue generated by those patents (by reference to the patent licensing table 1248).

Operators analyze the data obtained from the above processing to evaluate the customer's business from a patent/financial point of view.

It is noted that the steps of flowchart 10702 are provided for purposes of illustration, and not limitation. Additional steps directed to other financial operations that process the data in databases 316 can be added to the flowchart 10702. Additionally, the financial analysis process of FIG. 107 is an iterative one, as represented by control flow arrow 10712. Further, the steps of the financial analysis flowchart 10702 can be repeatedly performed for the same or different BOM groups, and can be performed in any order. Therefore, the steps of the financial analysis flowchart 10702 are provided in FIG. 107 for purposes of illustration, and not limitation.

Strategic Planning

The customer corporate entity can utilize the present invention to perform a strategic planning process. The overall goal of the strategic planning process is to develop a future product direction and/or future business strategy that cannot be easily duplicated by competitors. This future product direction and/or business strategy leverages the customer corporate entity's patent protection and technology ownership. The strategic planning process may result in considerable business process re-engineering based on future business analysis.

-205-

An example of the strategic planning process is represented by flowchart 10802 shown in FIG. 108. In step 10806, BOM data for a contemplated new product is received. This BOM data may have been generated by the R&D department, for example. This BOM data is loaded into the BOM databases 626 in the manner described above.

In step 10808, an operator associated with the customer corporate entity maps the customer corporate entity's patents (both owned and licensed) to the new future product. The operator also maps competitors' patents to the new future product.

In step 10810, the operator issues instructions to run patent mapping reports pertaining to the BOM group of the proposed future product. The operator also issues reports to run clustering/bracketing reports for the proposed future product. The patent mapping module 1002 and the patent clustering/bracketing module 1008 in the enterprise server 314 automatically runs these reports in the manner discussed above.

From these methodology reports, the operator is able to determine if the proposed future product is adequately protected by the patents which the customer corporate entity currently owns or licenses. If the proposed future product is not adequately covered by the customer corporate entity's current patents, then the customer corporate entity may consider whether or not it should prepare and file additional patent applications to further cover the proposed future product.

Also from these methodology reports, the operator can determine if the features of the proposed future product are covered by the patents owned by others. If some features of the proposed future product are covered by patents owned by others, then the customer corporate entity can consider whether it should license the patents, purchase the patents, modify the proposed future product in order to design around the patents, ignore the patents, abandon plans for the proposed future product, etc.

-206-

The above operation is described as being performed by a single operator. In practice, the above operation may require the involvement of many people with knowledge in many different fields and having different abilities (or expert systems with this knowledge and abilities).

5 The above analysis and decision making process performed by the customer corporate entity may involve running additional methodology reports, such as patent citation reports, patent aging reports, inventorship reports, and financial reports.

10 It is noted that the steps of the strategic planning flowchart 10802 shown in FIG. 108 are only representative of the type of tactical and strategic functions that can be performed based on the analysis and methodology reports that are automatically prepared by the analysis modules 416.

15 The strategic planning process is an iterative one, as represented by control flow arrow 10812. The steps of the strategic planning flowchart 10802 can be performed repeatedly, and in any order. Additional steps can be added to represent the additional strategic planning functions mentioned above. Accordingly, the steps of the strategic planning flowchart 10802 shown in FIG. 108 are provided for purposes of illustration, and not limitation.

Integrated Methodology Embodiment

20 The functions and methodologies described above can be performed individually or in any combination to achieve the objectives of the customer corporate entity. For example, at times, the customer may wish to run individual reports to examine a particular aspect of its business, such as the remaining patent terms of patents that map to a particular product line. At
25 other times, the customer may wish to run a combination of reports to analyze a larger part of its business, such as a strategic planning analysis to develop a patent strategy for a proposed future product. At still other times, the customer

-207-

may wish to perform an extensive and integrated examination of its business from a patent-centric and group-oriented point of view.

An example of such an integrated methodology is represented by a flowchart 10902 in FIG. 109. In this example, the patent mapping and mining process (described above with reference to FIG. 102) is combined with the
5 situation assessment process (see FIGS. 103-107) and the strategic planning process (see FIG. 108).

It should be understood that the patent mapping and mining process, the situation assessment process, and the strategic planning process can be performed
10 individually or in any combination with any other functions described herein according to the needs, requirements, and characteristics of the particular corporate entity being served.

User Interface

Referring to FIG. 114, the user interface module 11404 in the client 304,
15 306 supports a variety of user interface display formats, windows, and screens for enabling operators working at the client 304, 306 to view, manipulate, and otherwise interact with data. Many of the user interface display formats and display screens supported and provided by the user interface module 11404 are described above. Such user interface display formats and display screens
20 described above include: example user interface display formats pertinent to the searching features of the present invention shown in FIGS. 53-57, an example user interface display format pertinent to the display of group information shown in FIG. 58, examples of patent mapping display formats shown in FIGS. 59 and 60, examples of patent citation report display formats shown in FIGS. 61-65,
25 examples of patent aging display formats shown in FIGS. 66-70, examples of patent clustering/bracketing display formats shown in FIGS. 71-73, examples of inventor patent count display formats shown in FIGS. 74-77, examples of

-208-

employment information display formats shown FIGS. 78-80, an example display format showing the display of patent text in a first window and patent notes in a second window shown in FIG. 111, and an example display format showing the display of patent text in a first window and patent images in a second window shown in FIG. 112.

Other user interface display formats, display windows, and display screens supported and provided by the user interface module 11404 in the client 304, 306 shall now be described. It should be understood that the present invention is intended and adapted to support and provide display formats, display windows, and display screens other than those shown and described herein. In fact, the present invention can support and provide any display format, display window, or display screen useful for enabling operators to view and interact with the data and the data processing described herein. Accordingly, the display screens, display windows, and the display formats shown in the figures and discussed herein are provided solely for purposes of illustration, and not limitation.

User Login

FIG. 116 illustrates an example user login window 11602. An operator at a client 304, 306 enters information into the fields of the login window 11602 in order to log into and establish a new session with the Enterprise server 314.

In order to establish a session with the Enterprise server 314, the operator enters his user name in a login field 11608, and enters his password into a password field 11610. If the operator wishes to connect with the Enterprise server 314 via a specific server, then the operator enters the address of this server in a server field 11604. If the customer's network is configured to use "fire walls" with proxy servers, the operator enters information about the proxy server(s) in the proxy field 11606.

Console

FIGS. 18 and 58 depict example display formats wherein a group hierarchy is graphically illustrated in a first window or pane, and a list of patents and other documents contained in one of the groups of the group hierarchy is displayed in a second window or pane. These example display formats of FIGS. 18 and 58 are herein called the console display format, or simply the console.

FIG. 117 illustrates a console according to an alternative embodiment of the invention. The console 11702 includes a first window or pane 11704, a second window or pane 11706, and a third window or pane 11708. The first pane 11704 is also called the group pane 11704, the second pane 11706 is also called the document pane 11706, and the third pane 11708 is also called the notes pane 11708.

A group hierarchy is depicted in the group pane 11704. In the example of FIG. 117, the top level or root group in the group hierarchy is called the repository group 11710. The child groups of the repository group 11710 are not shown in FIG. 117 (i.e., the operator has not expanded the repository group 11710 in the group pane 11704). The child groups of the repository group 11710 are shown in FIG. 133 (i.e., the operator has expanded the repository group 11710 in the group pane 11704 in the example of FIG. 133).

Referring again to FIG. 117, the document pane 11706 includes a list of patents and other documents which are contained within a group selected from the group hierarchy depicted in the group pane 11704. The patents and documents are listed in a tabular or "spreadsheet" format. The list of patents and documents in the document pane 11706 includes both the patent numbers and patent bibliographical information for the patents, and bibliographic information for the non-patent documents. Such patent bibliographic information displayed in the document pane 11706 includes the title, abstract, inventor, class, issue date, and user-defined keywords. All additional patent bibliographic information can be

-210-

viewed in the document pane 11706 by utilizing the horizontal scroll bar 11720 to sideways scroll in the document pane 11706. Other embodiments of the invention allow the user to select an arbitrary number of bibliographic fields to view. In example of FIG. 117, no patents are listed in the document pane 11706 because a group has not been selected in the group hierarchy depicted in the group pane 11704.

The operator can restructure the group hierarchy displayed in the Group pane 11704 using well-known drag-and-drop operations. The operator can add any of the documents displayed in the Document pane 11706 to any group in the group hierarchy displayed in the Group pane 11704 by using well-known drag-and-drop operations.

The notes pane 11708 displays a list of the notes associated with either a group selected in the group pane 11704, or a patent or document selected in the document pane 11706. The list of notes in the notes pane 11708 is presented in a tabular or "spreadsheet" format. The list of notes in the notes pane 11708 includes information that identifies the type of the note (that is, either a patent/document note or a group note), and the title of the note. All other bibliographic or other information relating to notes can be viewed by manipulating the horizontal scroll bar 11722 in order to sideways scroll in the notes pane 11708.

As indicated above, FIG. 133 illustrates an example console 13302 wherein the repository group 11710 is expanded in the group pane 11704 so that its child groups are shown. In the example of FIG. 133, a group called "analog special effects controller subassembly" 13304 has been selected. Accordingly, a list of all patents and other documents in the selected group 13304 is displayed in the document pane 11706. In the example of FIG. 133, the selected group 13304 includes a single patent, U.S. Patent No. 5,029,013.

According to the present invention, the operator can view the text and/or image of any patent or other document listed in the document pane 11706 by

-211-

selecting the patent or document using well know item selection techniques. An operator can select a patent, for example, by double clicking on the patent number, or by using any other type of well known operator navigation procedure.

Suppose that the operator selects, from the document pane 11706, U.S. Patent No. 5,029,013. In accordance with this selection, the user interface module 11404 displays the text of the selected patent in a text window 12302, (see FIG. 123), and the image of the selected patent in an image window 12402 (see FIG. 124). The operator may selectively view the text window 12302 individually, the image window 12402 individually, or the text window 12302 and the image window 12402 simultaneously.

FIG. 118 illustrates an example console 11802 wherein the group hierarchy 11712 (with the repository group 11710 as the root) is depicted in the group pane 11704. The repository group 11710 is selected in the example of FIG. 118. Accordingly, a list of all patents and other documents in the repository group 11710 is displayed in the document pane 11706. Accordingly, by opening the folder corresponding to the repository group 11710 in the group pane 11704, the operator can view all patents that are contained in the patent repository (that is, all patents that are in the patent database 614), and, in fact, can view all documents that are in the repository.

Console Tool Bars

Referring again to FIG. 117, a Group toolbar 11724 is displayed in the group pane 11704. The tools in the Group toolbar 11724 are used to manipulate the groups of the group hierarchy displayed in the group pane 11704. A Document toolbar 11726 is displayed in the second pane. The tools in the Document toolbar 11726 are used to manipulate the documents listed in the document pane 11706. A Note toolbar 11728 is displayed in a notes pane 11708.

-212-

The tools in the Note toolbar 11728 are used to manipulate the notes listed in the notes pane 11708.

The Group toolbar 11724 is illustrated in greater detail in FIG. 135. The Document toolbar 11726 is illustrated in greater detail in FIG. 136. The Note toolbar 11728 is illustrated in greater detail in FIG. 137.

Referring to FIG. 135, an operator presses a Search button 13502 in the Group toolbar 11724 in order to conduct a search for data in the databases 316. An operator presses a Delete Group button 13504 in order to delete the group selected in the Group pane 11704. An operator presses a create New Group button 13506 in order to create a new group. The new group is created as a child group of the group selected in the Group pane 11704.

An operator presses a Show Group Properties button 13508 in the Group toolbar 11724 in order to view and edit the properties of the group selected in the Group pane 11704. An operator presses an Import Data button 13510 in order to import data from an external data file into the repository 612, 602. An operator presses an Export Data button 13512 in order to export data from the repository 612, 602 into a user specified location, such as an external file.

An operator presses a Refresh Data button 13514 in the Group toolbar 11724 in order to refresh the listing of all data in the Group pane 11704, the Document pane 11706, and the Notes pane 11708. Refresh is needed to synchronize multi-user access to the database. Such data refresh is performed by reading from the databases 316. An operator presses a Reports button 13516 in order to execute methodology reports.

Referring now to FIG. 136, an operator presses a Show Text button 13602 in the Document toolbar 11726 in order to display the text of a document selected in the Document window 11706. An operator presses a Show Image button 13604 in order to display the image of the document selected in the Document pane 11706. Examples of the display of text and images are shown in FIGS. 123 and 124.

-213-

An operator selects a Skim Images button 13606 in the Document toolbar 11726 in order to perform a Skim Images function. In accordance with the Skim Images function, the first image page of each document listed in the Document window 11706 is displayed in succession by repeatedly clicking the Show Images button 13606.

An operator presses an Add Patent button 13608 in order to add a new patent or document to the list of patents displayed in the Document pane 11706 (that is, to add a new patent or document to the group selected in the Group pane 11704). An operator presses a delete patent button 13610 in order to delete the patent or document selected in the Document pane 11706 from the group selected in the Group pane 11704 (this does not delete the document from the repository).

Referring now to FIG. 137, an operator presses an Add new Group Note button 13702 in order to add a new group note to a group selected in the Group pane 11704. An operator presses an Add New Patent/Document Note button 13704 in order to add a new patent note to the patent or document selected in the document pane 11706.

An operator presses a Delete The Selected Group Or Patent Note button 13706 in order to delete the note selected in the Notes window 11708. An operator presses the Sort Notes By Type button 13708 in order to sort the list of notes depicted in the Notes window 11708 by their type (that is, by group or by patent). An operator presses a Sort Notes By Title button 13710 in order to sort the list of notes depicted in the Notes window 11708 according to their titles.

Various ones of these functions that are invoked by pressing buttons in the Group toolbar 11724, the Document toolbar 11726, and the note toolbar 11728 are further described in the following sections.

-214-

Creating a New Group

An operator presses the Create New Group button 13506 in the Group toolbar 11724 in order to create a new group in the group hierarchy displayed in the Group pane 11704. The new group is created as a child group of the group selected in the Group pane 11704.

Consider, for example, the Example console 11902 illustrated in FIG. 119. Assume that the operator presses the Create New Group button 13506 when a group entitled "analog special effects controller subassembly" 11906 is selected. In this case, a new group 11904 is created as a child group of the selected group 11906. The default name of the new group 11904 is "A New Group". The name of the new group 11904 can then be modified using well-known procedures (such as by selecting the new group 11906, selecting a Rename command, and then typing in a new name for the new group 11904). FIG. 120 illustrates a scenario where the new group 11904 has been renamed as "Sample New Name."

Editing Group Properties

An operator presses the Show Group Properties button 13508 in the Group toolbar 11724 in order to view and edit the properties of the group selected in the Group pane 11704. Group Viewing And Editing windows 12602 are shown in FIGS. 126 and 127. Referring first to FIG. 127, a Group Properties tab 12606 is displayed. The operator can view and edit general information about the group selected in the Group pane 11704 by using the General Group Properties tab 12606. In particular, the operator can view and edit the group name via a Name field 12608, and a description about the group via a Description field 12610. The General Group Properties tab 12606 also displays the Group Identifier of the group selected in the Group pane 11704 in a field 12612.

-215-

FIG. 126 depicts a Group Properties Security tab 12604 in which the operator can view and enter security access information for the selected group. In particular, the operator can edit the users who have access to the selected group in a field 12608, and also can specify the access level for each of these users via fields 12610.

It is noted that access to the General Group Properties tab 12606 and the Group Properties Security tab 12604 is limited to operators having sufficiently high access levels with respect to the selected group.

Shared Groups

In an embodiment of the invention, shared groups are graphically depicted in the console using a stacked folders icon 15302 (FIG. 153). FIG. 154 shows an example console 15402 where both an "Only Luke can see!" group 15404 and an "Oscar Project" group 15408 share a "November Project" group 15406, which is a shared group.

FIG. 156 illustrates a group links tab 15602, which is the interface to additional group properties. These group properties in the group links table 15602 pertain to the links associated with a selected group. The scenario shown in FIG. 156 corresponds to the example of FIG. 154. Accordingly, in FIG. 156, the November Project group is shown as being a child of the Oscar project group and the Only Luke can see! group, which are both shown as being children of the Repository group. The operator can double click (or otherwise select) any link to open the lowest level data object, in this case the November Project group.

-216-

Invoking Patent-Centric and Group-Oriented Analysis Functions

An operator presses the Run Reports button 13516 in the Group toolbar 11724 in order to execute methodology reports. These methodology reports are performed by the Analysis modules 416 in the Enterprise server 316 (see FIG. 10), and are described above.

Upon the pressing of the Run Report button 13516, a Report Generator screen 12802 is displayed (FIG. 128). The operator selects the report or function that is to be executed by pressing a button 12806. This results in displaying a list of all available reports/functions. An example display of this function/reports list 12906 is shown in FIG. 129. The operator selects one of the function/reports from this list in a well-known manner. In the example of FIG. 128, the operator has selected the inventor employment information function (which is executed by the inventor employment information module 1014 of the Enterprise server 314).

The operator can enter information into the other fields of the report generator screen 12802 in order to define and/or limit the scope of the function that is to be performed. In particular, the operator can specify the number of levels in the group hierarchy to consider in performing the function (specified in field 12804) by entering the appropriate level information into a field 12810. If, for example, the operator enters the value 1 into field 12810, then only the group selected in the Group pane 11704 is considered in performing the selected function. If, instead, the operator enters the value 2 into field 12810, then both the selected group and its immediate child group (one level down from the selected group) are considered in performing the selected function. In the example of FIG. 128, the operator has entered the value 5 into field 12810. Accordingly, in performing the inventor employment information function, both the selected group and four levels of its child groups are considered.

The operator can limit the patents which are to be considered in performing the selected function by patent expiration dates by entering information

-217-

into fields 12812. In the example of FIG. 128, all patents which are expiring or which have expired between the dates of January 1, 1900 and December 31, 2100 are considered in performing employment information function.

5 The operator can also limit the patents which are to be considered in performing the selected function according to the remaining patent terms of the patents. This information is entered into fields 12814. In the example of FIG. 128, only patents with remaining patent terms between three years and 10 years are considered in performing the inventor employment information function.

10 Once the operator has defined the function that is to be performed by entering information into the report generator screen 12802, the operator presses the Execute button 12808 in order to execute the selected function.

Adding Documents to a Group

15 The operator presses an Add a Patent/Document button 13608 in the Document toolbar 11726 in order to add a patent or some other document to the group selected in the Group window 11704. Upon pressing the Add a Patent/Document button 13608, a new Document window 13006 is displayed (FIG. 130). The operator inserts information identifying the document that is to be added to the selected group in a Document Name field 13008. In an alternative embodiment, the operator can select a document for insertion into the selected group by browsing through a graphical list of the documents in the Document databases 612, and use well-known drag and drop techniques to place documents into the selected group.

Adding a Document Note

25 The operator presses the Add a New Patent/Document Note button 13704 in the note toolbar 11728 in order to add a new note to a document. Consider the

-218-

example shown in FIG. 125. In this example, a Text window 12506 displays a portion of the text from U.S. Patent No. 5,029,013. The present invention allows the operator to attach a note to any portion of this patent. Suppose that the operator wishes to attach a note to the title, "Apparatus for Editing Videotapes" 12510. According to the present invention, the operator selects a pen 12508, and then uses the selected pen 12508 to highlight the title 12510. The operator then presses the Add a New Patent/Document Note button 13704. This results in a new note being created. This new note is displayed in a Note window 12514. The operator can enter a title for the note in field 12520, and can enter the text or other information for the note in field 12524. The operator can also indicate whether this note represents attorney work product via field 12522.

The Note window 12514 includes location information that identifies the location of the text in the document to which the note is linked. This location information is represented by reference numbers 12516 and 12544. The location information 12516 indicates that the note is attached to a portion of U.S. Patent No. 5,029,013 starting at line 6. The location information 12544 indicates that this portion in U.S. Patent No. 5,029,013 represents bibliographical information portion (indicated by the symbol B1).

In some embodiments, it is not necessary for the operator to press the Add a New Patent/Document Note button 13704. Instead, the new note is automatically created once the selected pen 12508 is used to highlight the title 12510. Other aspects of the present invention relating to creating and manipulating notes are described in U.S. Patent No. 5,623,681, U.S. Patent No. 5,623,679, pending U.S. Application Serial No. 08/341,129, and pending U.S. Application Serial No. 08/590,082, incorporated herein by reference in their entireties.

The Notes window 12514 includes an Unlink button 12534 which, when pressed, operates to unlink the note from the document. The note is then called a floating note. The Notes window 12514 also includes a Contract button 12536

-219-

which operates to reduce the area on the display screen which is used to display notes. The Notes window 12514 also includes an Expand button 12538 which operates to increase the amount of space on the display screen which is used to display notes.

5 As described above, the Notes pane 11708 in the console 12502 lists the notes associated with the selected document in the Document pane 11706, or the selected group in the Group pane 11704. In the example of FIG. 125, the Notes window 11708 lists the notes associated with the document selected in the Document pane 11706. The document selected in Document pane 11706 is U.S.
10 Patent No. 5,029,013. As just explained above, the operator has created a new note that is linked to this selected document. The title of this new note is "Simple Note". Accordingly, this note is listed in the Notes window 11708. The listing corresponding to this new note indicates the type of a note (in this case, the type is P, indicating a patent note) and the title of the new note (in this case, the title
15 is "Simple Note"). Additional information relating to the new note, and any other notes listed in the Notes pane 11708, can be viewed by sideways scrolling in the Notes pane 11708 by use of the horizontal scroll bar 11722.

Adding a Group Note

20 The operator presses the Add a New Group Note button 13702 in the Note toolbar 11728 in order to create a new group note for the group selected in the group pane 11704. Upon selection of the Add a New Group Note button 13702, a New Group window 13406 is displayed (FIG. 134). In the New Group window 13406, the user can enter the title of the group note in field 13410, and can enter information for the group note in field 13408. The operator can also
25 indicate whether the group note constitutes attorney work product by checking field 13412.

-220-

5 The Notes pane 11708 lists the notes corresponding to the document selected in the Document pane 11706, or the group selected in the Group pane 11704. In the example of FIG. 134, the group titled "Analog special effects controller subsystem" 13404 is selected. Accordingly, the Notes pane 11708 lists all group notes associated with this selected group 13404. The New Group note (represented by the data in the New Group window 13406) is linked to the selected group 13404. Accordingly, the new group note is listed in the notes pane 11708. The listing for this new group note indicates the type of the note (in this case, the type is C, for a case or group note) and the title of the group note (in the example of FIG. 134, the title is "Purpose of this group"). Other information pertaining to the group notes can be viewed in the Notes pane 11708 by horizontally scrolling in the Notes pane 11708 by using the horizontal scroll bar 11722.

Searching

15 The operator presses the Search button 13502 in the Group toolbar 11724 in order to perform a search of the data in the databases 316. Upon pressing the Search button 13502, a search screen 12102 is displayed (FIG. 121). Other example search screens are shown in FIGS. 53 and 57. The operator enters information into the Search screen 12102 in order to define the parameters of a new search. In the example of FIG. 121, the operator has entered the phrase "video tape" into an abstract field 12104. Accordingly, the operator has indicated that he wishes to conduct a new search for all documents where the phrase "video tape" appears in the abstract.

25 FIG. 122 illustrates an example console 12202 that results after the search defined in the example Search screen 12102 (FIG. 121) is performed. According to the present invention, a new temporary group 12204 (called New Search) is created to hold the search results (in other words, the documents that are

-221-

identified from performance of the search are placed in the new group 12204 called "New Search"). This new group 12204 is created as a child group of the group that was selected in the group pane 11704 when the search was executed.

A list of the documents identified by operation of the search is displayed
5 in the Document pane 11706. The operator can add these documents to any other group in the group hierarchy (displayed in the Group pane 11704) using well-known drag-and-drop operations. Alternatively, the operator can convert the new group 12204 to a permanent group (pre-defined or user-defined), and maintain these documents in the permanent group 12204 entitled "New Search".

10 FIG. 155 illustrates an example console 15502 where a temporary group 15504 called "New Search Group" has been created. Temporary groups may be color coded differently from other groups in the group hierarchy 15508. In the embodiment of FIG. 155, the group toolbar 11724 includes a "regular group" button 15506. By pressing the regular group button 15506, the temporary group
15 15504 becomes a permanent group. In an embodiment of the invention, the type of the new permanent group is the same as the permanent group's immediate parent. Thus, if the parent is a user-defined group, then the new permanent group is a user-defined group. If the parent is a BOM group, then the new permanent group is a BOM group. In an alternate embodiment, the operator is prompted to
20 enter a type for the temporary group 15504. In another embodiment, the type is predefined, such as BOM.

In an embodiment of the invention, a temporary group is automatically deleted at some time in the future, such as after a timeout or at the conclusion of the session with the enterprise server 314, if not converted to a permanent group
25 before then.

Web Searching

The operation of a web client 304 with regard to searching for data in the databases 316 according to an embodiment of the invention is described in this section. It should be understood, however, that the search capabilities described
5 herein are applicable to all clients 304,306 in various embodiments of the invention. That is, in some embodiments of the invention, all clients 304,306 include at least the capabilities described herein.

In the following, searching at the web clients 304 is described in the context of patent searching. However, the following description applies to
10 searching for all types of documents.

An operator at a web client 304 presses the Search button 13502 in the Group toolbar 11724 in order to access the search capabilities of the present invention. Upon pressing the Search button 13502, a Patent Search screen or window 14002 is displayed (FIG. 140). This Patent Search screen 14002 is
15 similar to the one shown in FIG. 53. In fact, some embodiments use the Patent Search screen 5302 shown in FIG. 53. Other embodiments use both the Patent Search screens 5302 and 14002 shown in FIGS. 53 and 140, respectively.

The operator enters information into the fields of the Patent Search screen 14002 in order to define the parameters of the search. For example, the operator
20 can define the search in terms of patent number, title, inventor, assignee, class, user-defined key words, date of issue, abstract, and/or full patent text by entering search terms into the corresponding fields of the Patent Search screen 140. Also, the operator can select which fields to display in the search results by appropriate selection of the Check Box fields 14004. Further, the operator can order the
25 display or printing of the search results according to a number of factors, such as patent number, assignee, expiration date, number of years remaining in patent term, or score. The score corresponds to the number of hits of the search parameters in a patent. The operator orders the search results by appropriate

-223-

selection in the Order field 14006. The operator can also specify the number of patents in the search results to display per screen by entering the appropriate information in field 14008.

5 In the example of FIG. 140, the operator has entered the search term "PCMCIA" in the Abstract field. In other words, the search defined by the operator is one that will identify all patents having abstracts with one or more occurrences of the word "PCMCIA". After the operator has fully defined the search in the Patent Search screen 14002, the operator presses the search button 14010 in order to execute the search.

10 FIGS. 141-143 illustrate example Search Results screens corresponding to the search specified in the example Patent Search screen 14002 of FIG. 140. The execution of this search identified 85 patents. This information is indicated at reference number 14104 in Search Results screen 14102. Search Results screen 14102 displays information on the first 10 of these 85 patents. This is indicated
15 at reference number 14106. Search Results screen 14202 (FIG. 142) displays information on the second 10 of these 85 patents. This is indicated at reference number 14204. Search Results screen 14302 displays information corresponding to the third 10 of these 85 patents. This is indicated at reference number 14304. The operator can display information on different ones of the 85 patents by use of
20 the navigation arrows 14108. In other words, the operator can scroll through search results screens 14102, 14202, 14302 by use of the navigation arrows 14108.

25 Referring to FIG. 141, the information on the patents identified by performing the search is presented in a tabular format. The information displayed for each patent includes the score (that is, the number of hits of the search parameters in the patent), the patent number, and the title. More or less information can be displayed for each patent. This is controlled by the operator by appropriately checking the check boxes 14004 in the Patent Search screen 14002.

-224-

The information for each patent also includes an indication of whether or not the patent is in the local repository (i.e., stored in the patent database 614). This is indicated in the column called "In Repository?" In the example of FIGS. 141-143, for example, U.S. Patent No. 5,334,030 is not in the patent database 614. However, U.S. Patent No. 5,540,597 is stored in the patent database 614.

The operator can view additional information on any of the patents listed in the Search Results screens 14102, 14202, 14302 by selecting the patent (for example, by double-clicking on the patent number in the list of patents). The extent of the information that is displayed for a selected patent depends on whether or not the patent is stored in the patent database 614. If the patent is not stored in the patent database, then patent bibliographic information (retrieved from the patent bibliographic databases 604) is displayed for the selected patent. An example display of such patent bibliographic information for a selected patent is shown in FIG. 144. The information in this example display 14402 corresponds to U.S. Patent No. 5,183,404 (which the operator selected by selecting entry 14110 in the Search Results screen 14102 of FIG. 141).

The display screen 14402 includes patent bibliographic information 14404 on U.S. Patent No. 5,183,404. Also, in accordance with an embodiment of the invention, the display screen 14402 also includes text corresponding to the abstract 14406 of U.S. Patent No. 5,183,404. In other words, for patents not stored in the patent database 614, the system of the present invention maintains both patent bibliographic information and the abstract. Such patent bibliographic information and abstract are stored in the patent bibliographic databases 604. (For example, the abstract can be stored in an Abstract database that is keyed to the primary patent table 1222 according to document_ID.)

FIGS. 145A-145C depict example patent information screens that display patent information on patents that are contained in the Patent database 614. The Patent display screen 14502 in FIG. 145A corresponds to U.S. Patent No.

-225-

5,615,328 (which the operator selected by double-clicking on entry 14310 in the Search Results screen 14302 of FIG. 143).

The operator can elect to display the text of the selected patent by pressing the Text button 14504. In the example of FIGS. 145A-145C, the operator has
5 pressed the Text button 14504. Accordingly, the text of the selected patent is displayed in the Patent Information screens of FIGS. 145A-145C. FIGS. 145A-145C represent the first three (of 15) electronic pages of the text of the selected patent.

The operator presses an Image button 14506 to display the images of the
10 selected patent. An example of such a display is shown in FIG. 56 (although the example in FIG. 56 does not correspond to U.S. Patent No. 5,615,328). The operator presses a Text And Image button 14508 to simultaneously display the text and images of the selected patent. An example display of such text and image is shown in FIG. 112 (although the example display in FIG. 112 does not
15 correspond to U.S. Patent No. 5,615,328 shown in FIG. 145A). If the operator wishes to initiate a new search, the operator presses a New Search button 14510. This causes the Patent Search screen 14002 (FIG. 140) to be displayed.

The selected patent as displayed in the display screens of FIGS. 145A-145C is internally hyperlinked. Such hyperlinks are represented by buttons 14524.
20 This button represented hyperlinking is a well known use of HTML.

If the operator wishes to view the bibliography of the selected patent, for example, then the operator presses the Bibliography button 14512. The operator presses the Cross Reference To Related Applications button 14516 to view cross-reference information of the selected patent. The operator presses the Brief
25 Description Of The Drawings button 14518 to view the brief description of the drawings section of the selected patent. The operator presses the Detailed Description Of The Invention button 14520 to view the detailed description section of the selected patent. The operator presses the Claims button 14522 to view the claims section of the selected patent.

Preferably, the invention implements such internal hyperlinking using a linked list of nodes, wherein each node corresponds to a section of the selected patent. Specifically, each node includes a link to the corresponding section of the selected patent. When the operator selects one of the Internal Linking buttons 14524, the invention traverses through the linked list corresponding to the selected patent until it finds the node corresponding to the section that the operator wishes to view. The invention then follows the link stored in the node to access the data corresponding to the section that the operator wishes to view. Additional details pertaining to this linked list is described in U.S. Patent No. 5,623,679, incorporated by reference herein.

The Patent Information display screens in FIGS. 145A-145C may also linked to external objects, such as other patents, documents, etc. In other words, the selected document is linked to external objects. For example, if the operator wishes to view the patent representing the immediate parent of the selected patent, then the operator presses a Parent Case Text button 14514. This results in the parent patent being retrieved and displayed.

The selected patent may also be linked to references to other documents that are contained in the selected patent. For example, the selected patent as displayed in FIGS. 145A-145C includes links to cited U.S. patent documents. These links are represented by reference numbers 14526 and 14528. Hyperlinks in the selected patent are denoted by underlining. Alternatively, any other method of highlighting can be used, such as fonts, color coding, symbols, icons, etc.

The selected patent may also include links to citations of other publications, as represented by reference number 14530, and citations to related applications, as represented by reference number 14532.

For reference purposes, citations to documents that are linked to the actual documents are called linked citations.

The operator can view the document corresponding to any linked citation by selecting that linked citation. For example, FIG. 146 illustrates a display screen

-227-

that shows information pertaining to U.S. Patent No. 5,206,830. This display screen 14602 was generated pursuant to the operator selecting a corresponding linked citation 14536 in the selected patent (see FIG. 145B). In the display screen 14602 of FIG. 146, the operator can elect to electronically order a copy of the patent from a third party provider by pressing an Order button 14604.

Referring again to FIGS. 145A-145C, not all document citations contained in the selected patent are linked to the cited documents. For example, referring to FIG. 145B, the citation to U.S. Patent No. 5,430,681 is not linked to the cited document (that is, it is not linked to an electronic copy of U.S. Patent No. 5,430,681). According to the present invention, a citation in a document being displayed is linked to the cited reference only if that cited reference is represented in the document bibliographic databases 602 and/or the document databases 612. In other words, the documents corresponding to linked citations are represented in the document bibliographic databases 602 and/or the document databases 612.

A citation in a document being displayed is not a linked citation if the corresponding document is not represented in the document bibliographic databases 602 or the document databases 612. Accordingly, in the example of FIG. 145B, the document bibliographic databases 602 and the document databases 612 include no information pertaining to U.S. Patent No. 5,430,681.

In other words, whether a citation to a document is linked to the document depends on whether or not the databases 316 store information on that document. If the databases 316 store information on the document, then the citation to the document is linked to the document (i.e., the citation is a linked citation). Otherwise, the citation to the document is not linked to the document (i.e., the citation is not a linked citation).

According to an embodiment of the invention, hyperlinking information is stored in a table in the databases 316. This hyperlinking table stores information pertaining to the links between documents. For a particular document, the

-228-

hyperlinking table stores information that identifies the links between that particular document and other documents (these other documents being cited in the particular document). More particularly, each entry of this table corresponds to a linked citation, and stores information identifying both ends of the link of the
5 linked citation. Specifically, each entry stores information that identifies a first document and the linked portion in the first document, and also includes information that identifies the object to which the linked portion in the first document is linked. This object may be a second document, for example. In the example of FIGS. 145A-145C, the hyperlinking table includes -hyperlinking
10 information that identifies the documents that are linked to the patent being displayed (that is, U.S. Patent No. 5,615,328).

In an alternate embodiment, when a document is displayed, the document is automatically analyzed to identify and generate a list of all citations to other documents contained in the document. This list is then compared to the
15 information stored in the databases 316. If the databases 316 have stored therein information on a document in the list, then the citation to that document becomes a linked citation. Otherwise, the citation to the document does not become a linked citation.

The invention also allows a document to have user-defined links. In
20 operation, an operator would select a portion of a document being displayed, and then link that portion to some object. That object may be internal to the document (that is, linking a portion of a patent to another portion of the same patent), or external to the document (for example, linking a portion of a patent to another document or to a website). One portion corresponding to a user-defined link can
25 completely or partially overlap with another portion corresponding to a user-defined link.

For example, as represented in FIG. 145C, the user has selected the company name "International Business Machines Corporation" and has linked this selected portion 14534 of the document being displayed to an external object. In

the example of FIG. 145C, the operator has linked this portion 14534 of the document being displayed to the IBM patent website. Accordingly, when the operator selects this linked portion 14534, the IBM website is accessed.

Preferably, information pertaining to user-defined links is stored in a user-defined linking table in the databases 316. Each entry of this table corresponds to a user-defined link, and stores information identifying both ends of the link. Specifically, each entry stores information that identifies the document and the linked portion in the document, and also includes information that identifies the object to which the linked portion in the document is linked.

In practice, when the operator selects a linked citation, the client 304,306 generates a request to the Enterprise server 314 to retrieve information pertaining to the document corresponding to that linked citation (assuming that such information is not already stored in the cache within the client 304,306). In the case of web clients 304, patent-centric URL commands are sent to the Enterprise server 314 (see FIG. 151).

Referring again to the Search Results screens in FIGS. 141-143, the operator can press a Get Results In File button 14194 to write the search results to a user-specified file. The operator can also press an "In Repository" button 14190 to view a subset of the search results corresponding to patents that are in the local patent repository (that is, patents that are in the Patent database 614). An example display screen that results from pressing the "In Repository" button 14190 is presented in FIG. 147.

Still referring to FIG. 141, the operator can press a "Not In Repository" button 14192 in order to view a subset of the search results corresponding to patents that are not in the Patent database 614. This list of patents not in the Patent database 614 may be useful to the operator. For example, the operator can generate a purchase order to obtain these patents using this list.

Referring now to FIG. 147, the Search Results: Patents In Repository display screen 14702 includes a Skim Images button 14704. The operator presses

-230-

the Skim Images button 14704 in order to view the first image page of each of the patents listed in table 14706 (this table lists the patents from the search results that are contained in the Patent database 614). The operator displays the first image page of successive patents listed in table 14706 by repeatedly clicking the Skip Images button 14704.

A Skip Images display screen 14802 is shown in FIG. 148. This Skim Images Display screen 14802 results from clicking the Skip Images button 14704 a first time (that is, the image displayed in the Skip Images display screen 14802 corresponds to the first page of the first patent listed in table 14706). The operator can view the first image page of either a previous patent or a next patent in the list 14706 by use of the Navigation buttons 14806.

With regard to the web client 304, pressing the Skip Images button 14704 causes the web browser 808 to generate a patent-centric URL command. This patent-centric URL command represents a request to the Enterprise server 314 to retrieve the first image page of the associated patent (that is, the next patent in the list 14706). Enterprise server 314 responds by returning raw data corresponding to this image page. The translator 804 in the web server 310 converts this raw data to HTML data. The HTML data is sent to the web client 304. The browser 808 in the web client 304 receives this HTML data, and renders the HTML data in order to display the first image page of the patent, as shown by way of example in FIG. 148.

An example patent-centric URL command is displayed in field 14804 of the Skim Images display 14802.

The Skim Images feature of the present invention is analogous to a manual search of patents in the shoes of the USPTO Search Room. During such manual searches, practitioners often quickly thumb through the patents in a patent shoe by looking at the first pages of the patents. The Skim Images feature of the present invention enables an operator to quickly electronically scan over the first image pages of the patents in the list 14706, thereby emulating a manual search

-231-

through the shoes in the PTO Search Room. It should be noted that the images displayed to the user are preferably HTML data rendered by the browser 808.

Importing Data

The operator presses the Import Data button 13510 in the Group toolbar 11724 in order to import data from an external location into the repository 612,602. Upon pressing the Import Data button 13510, an Import Data window 13102 is displayed (FIG. 131). In the Import Data window 13102, the user specifies the external drive in field 13104 (the user can browse through the available storage locations by pressing button 13106). The files in this selected location are displayed in window 13108. The operator enters the name of the file to import in field 13110. The type of this file (such as text, database, spreadsheet, etc.) is entered in field 13112. After specifying the parameters of the import operation, the operator presses the Open button 13190 in order to import the selected file from the selected external storage location into the repository 612, 602.

Exporting Data

The operator presses the Export Data button 13512 in the Group toolbar 11724 in order to export data from the repository 612,602 to an external location, such as an external file. Upon pressing the Export Data button 13512, an Export window 13202 is displayed (FIG. 132). In the Export window 13202, the operator enters the external location of the file to which the export operation is to be directed. This external location is entered in field 13204. It is noted that the operator can browse through all available external locations by pressing button 13206. A list of the files located in the external location (specified in field 13204) is displayed in field 13208. The user enters the name of the external location (to

-232-

which the data is to be exported) in field 13210. The type of the file (such as text, spreadsheet, database) is entered in field 13212. After the operator has specified the parameters of the export operation, the operator presses the Save button 13290 in order to execute the export data operation.

5 *Data Presenting and Processing Using Hyperbolic Trees*

As described above, hyperbolic trees can be used for processing and presenting data in the context of the patent citation function. In the following, example uses of hyperbolic trees for data presentation and processing in the context of the present invention are described. First, the hyperbolic tree mechanism is generally described. Next, various functions of the present invention that optionally use hyperbolic trees for data presentation and processing are described.

General Description of Hyperbolic Trees

A hyperbolic tree or hyperbolic browser is a well known "focus + context" technique for visualizing and manipulating large hierarchies, such as trees. A hyperbolic browser assigns more display space to a portion of the tree while still embedding it in the context of the entire tree. The essence of this scheme is to lay out the tree in a uniform way on a hyperbolic plane and map this plane onto a circular display region. Thus, hyperbolic trees support a smooth blending between focus and context, as well as continuous redirection of the focus. Well known procedures exist for manipulating the focus using pointer clicks as well as interactive dragging, and for smoothly animating transitions across such manipulation.

-233-

A hyperbolic browser initially displays a tree with its root at the center, but the display can be smoothly transformed to bring other nodes into focus. This is shown in the example of FIGS. 177 and 178. FIG. 177 illustrates an example hyperbolic tree 17701. A portion of the hyperbolic tree 17701 is displayed in a window 17702. The hyperbolic tree 17701 includes a root node 17704 and a plurality of additional nodes connected either directly or indirectly to the root node 17704. For example, node 17706 is connected directly to root node 17704. Node 17708 is connected indirectly to root node 17704.

Additional portions of the hyperbolic tree 17701 can be expanded and displayed in the window 17702 (that is, focus can be redirected to those portions) by traversing to the desired portions in the hyperbolic tree 17701. An operator traverses the hyperbolic tree 17701 using a pointer device, such as a mouse. For example, suppose that the operator wishes to focus the hyperbolic tree 17701 at node 17708. In this case, the operator moves the mouse (represented by icon 17703) and clicks proximate to node 17708. In response, the portion of the hyperbolic tree 17701 proximate to node 17708 is expanded and displayed. This is shown in FIG. 178.

With hyperbolic trees, the amount of space available to a node falls off as a continuous function of its distance in the tree from the point in the center. Thus, the context always includes several generations of parents, siblings, and children, making it easier for the user to explore the hierarchy without getting lost.

Hyperbolic trees exploit well known hyperbolic geometry. Hyperbolic geometry is described in a number of publicly available documents, such as Coxeter, *Non-Euclidean Geometry*, University of Toronto Press, 1965, and Moise, *Elemental Geometry from an Advanced Standpoint*, Addison-Wesley, 1974. As is well known, the essence of hyperbolic trees is to lay out the hierarchy, such as a tree, on a hyperbolic plane and map this plane onto a circular display region. The hyperbolic plane is a non-Euclidean geometry in which parallel lines diverge away from each other. This leads to the convenient property that the

-234-

circumference of a circle on the hyperbolic plane grows exponentially with its radius, which means that exponentially more space is available with increasing distance. Thus, trees -- which tend to expand exponentially with depth -- can be laid out in hyperbolic space in a uniform way, so that the distance (as measured
5 in the hyperbolic geometry) between parents, children, and siblings is approximately the same everywhere in the hierarchy.

While the hyperbolic plane is a mathematical object, it can be mapped in a natural way onto the unit disk, which provides a means for displaying it on an ordinary (Euclidean) display. This mapping displays portions of the plane near the
10 origin using more space than other portions of the plane. Very remote parts of the hyperbolic plane get minuscule amounts of space near the edge of the disk. Translating the tree on the hyperbolic plane provides a mechanism for controlling which portion of the structure receives the most space without compromising the illusion of viewing the entire hyperbolic plane.

15 Much of the above general description of hyperbolic trees and browsers was taken or adapted from Lamping *et al.*, "A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies," *CHI '95 Proceedings Papers*, Xerox Palo Alto Research Center, 1995 (including the Appendix entitled "Implementing Hyperbolic Geometry"), incorporated herein by reference in its
20 entirety. In addition to conventional sources, the Lamping *et al.* document and its Appendix are available on the Internet at www.acm.org/sigchi/chi95/Electronic/documnts/papers/jl_bdy.htm. Hyperbolic trees and browsers are described in a number of additional publicly available documents, such as *Hyperbolic Tree Toolkit Application: Programmer's Reference*, Version 1.0, Xerox Corporation and InXight Software, Inc., January
25 1997, which is incorporated herein by reference in its entirety.

Each node in a hyperbolic tree, when being displayed, includes one parent node and zero or more child nodes. The one exception is the root node, which does not have a parent node, but which has zero or more child nodes. For

-235-

example, node 17706 in FIG. 177 includes one parent node (node 17704) and three child nodes (nodes 17750, 17752, and 17754).

This description of the hyperbolic tree is a general description of the "tree" construct. That is, a tree includes a root node and multiple non-root nodes each having a single parent node and zero or more child nodes.

A graph, such as a directed acyclic graph (DAG), differs from a tree. Like a tree, a DAG includes a root node that does not have any parent nodes, and that has zero or more child nodes. (Generally, a DAG is not limited to a single root, or parentless, node. However, DAGs when used with the present invention either include a single root node, or are made to include a single root node. This is done by creating an artificial root node, and adding to it as children all of the previously parentless nodes.)

Also like a tree, each non-root node in a DAG includes zero or more child nodes. However, unlike trees, each non-root node in a DAG includes one or more parent nodes. An example DAG 17902 is shown in FIG. 179 where non-root node 17906 includes parent nodes B and C.

Trees and graphs (such as DAGs) are well known and are described in many publicly available documents, such as Aho *et al.*, *Data Structures and Algorithms*, ISBN 0-201-00023-7, 1983.

Many of the modules of the present invention generate data which is optimally represented as a graph, such as a DAG. For example, patent citation data is best represented as a DAG since any given patent can be cited in multiple patents, and can also itself cite multiple patents. Also, claim charts are optimally represented as DAGs since any given dependent claim can depend from multiple parent claims (that is, a multiple dependent claim).

At the same time, the data generated by modules of the present invention can often be effectively represented through use of hyperbolic trees. This is the case, for example, with multi-level patent citation data (as described above).

-236-

Accordingly, various modules of the present invention enable the operator to display data using hyperbolic trees.

Preferably, the present invention represents data, where appropriate, as a graph, such as a DAG. When an operator elects to display the data using a hyperbolic tree format, the invention does so.

Some well known hyperbolic browser implementations are capable of generating and displaying hyperbolic trees directly from DAGs. These hyperbolic browser implementations conceptually operate by mapping the root node in the DAG to the root node in the hyperbolic tree. Also, non-root nodes in DAG that have a single parent node are directly mapped to corresponding nodes in the hyperbolic tree. These hyperbolic browser implementations conceptually map non-root nodes in the DAG that have multiple parent nodes to multiple nodes in the hyperbolic tree, wherein each of these tree nodes are linked to a single parent node in the hyperbolic tree. This is shown in FIG. 179, where node D has parent nodes B and C in DAG 17902. As a result of this conceptual mapping process in hyperbolic browser, nodes 17906A and 17906B are created in the hyperbolic tree 17904. Each of these nodes 17906A and 17906B is linked to a single parent node (that is, nodes B and C) in the hyperbolic tree 17904. The process of generating and populating a hyperbolic tree from a DAG by such hyperbolic browser implementations will be apparent to persons skilled in the relevant art(s).

When operating with such hyperbolic browser implementations, the invention passes the DAG to the hyperbolic browser. The hyperbolic browser generates and displays a hyperbolic tree directly from the DAG.

Other well known hyperbolic browser implementations are not capable of generating and displaying hyperbolic trees directly from DAGs. When working with such hyperbolic browser implementations, the invention translates or maps the DAG to a tree format. This mapping process is illustrated in FIG. 179. According to this mapping process, the root node in the DAG 17902 is directly mapped to the root node in the tree 17904 (see node A). Also, non-root nodes

-237-

in DAG 17902 that have a single parent node are directly mapped to corresponding nodes in the tree 17904 (see nodes B and C). Non-root nodes in the DAG 17902 that have multiple parent nodes are mapped to multiple nodes in the tree 17904. Each of these tree nodes are linked to a single parent node in the tree 17904. See node D, which has parent nodes B and C in DAG 17902. As a result of the mapping process, nodes 17906A and 17906B are created in the tree 17904. Each of these nodes 17906A and 17906B is linked to a single parent node (that is, nodes B and C).

The invention then passes the tree 17904 to the hyperbolic browser. The hyperbolic browser generates and displays a hyperbolic tree from the tree 17904. The process of generating and populating a hyperbolic tree from the tree 17904 will be apparent to persons skilled in the relevant art(s).

In some embodiments, the invention performs this mapping process even when working with a DAG-enabled hyperbolic browser. This may be done for performance reasons. For example, performance may be enhanced by having the invention perform the mapping function, as opposed to the hyperbolic browser performing the mapping function.

Since the invention generates a graph or DAG 17902 to represent data, where appropriate, it is possible to easily display the information in DAG 17902 using other display mechanisms that are DAG-capable or DAG-enabled, or more generally graph-capable or graph-enabled. The scope and spirit of the present invention is intended and adapted to operate with DAG-capable or DAG-enabled display mechanisms, or more generally graph-capable or graph-enabled mechanisms, including those presently existing and those that will be developed in the future.

Various functions of the present invention that optionally use hyperbolic trees to present and process data are described in the following sections.

Patent Citation Tree

According to an embodiment of the present invention, the patent citation module 1004 performs a patent citation tree function (in addition to the functions described above).

5 FIG. 164 illustrates an example patent citation tree 16404, a portion of which is being displayed in window 16412. In the patent citation tree 16404, each node represents a patent, and each link represents a citation. When performing a forward patent citation function, the links going from the root node toward the children nodes or leaves represent the directed association "is cited in," as in
10 "<parent-node> is cited in <child-node>." When performing a backward citation function, the links going from the root node toward the leaves represent the directed association "cites," as in "<parent-node> cites <child-node>."

 As discussed below, the nodes in the patent citation tree 16404 are preferably color coded or otherwise emphasized (by modifying the display
15 properties of the nodes) according to patent ownership or assignee, or some other criteria, such as time (i.e., issue date) or user-selection (this is described below). As a result, the patent citation tree 16404 is very useful when performing clustering analysis and/or bracketing analysis. In fact, the patent clustering/bracketing module 1008 in some embodiments is capable of using
20 hyperbolic trees to represent its data output. Accordingly, the description below of the patent citation module 1004 when performing the patent citation tree function using hyperbolic trees is also applicable to the patent clustering/bracketing module 1008. Additionally, other modules of the invention are capable of presenting their data output using hyperbolic trees. Accordingly,
25 the following description is also applicable to these other modules of the invention.

-239-

Patent Citation Tree (Network Client)

The patent citation tree function when being performed by a network client 306 interacting with the enterprise server 314 shall now be described with reference to a flowchart 15702 in FIG. 157. For illustrative purposes, suppose that an operator at a network client 306 is viewing the example console 16102 shown in FIG. 161.

In step 15706, the operator selects a patent from a listing of patents displayed in document pane 16104. The patent selected shall be referred to as the '484 patent. The operator selects the '484 patent using a pointing device, such as a mouse (represented by icon 16110). An indication of this selection of the '484 patent is received by the network client 306.

In step 15708, while the '484 patent is selected, the operator issues a command to perform a citation analysis function. The operator may issue this command by right-clicking (while the mouse is positioned over the '484 patent in the document pane 16104) to bring up a context-sensitive pop-up menu 16202, a portion of which is shown in FIG. 162. Alternatively, the operator can press a button associated with the citation analysis function in the toolbar 16108 (which has been modified to include options in addition to those described above, such as the citation analysis function option), or make a selection from a menu accessible via the menu bar. The network client 306 receives an indication of this command issued by the operator.

In step 15710, the network client 306 presents a dialogue box 16302 (FIG. 163) to the operator. In this dialogue box 16302, the operator identifies whether he wishes to perform a forward citation function or a backward citation function. Also, the operator defines the scope of the citation analysis to be performed by indicating the number of levels that should be cited. This information entered into the dialogue box 16302 is received by the network client 306.

In step 15712, the network client 306 interacts with the enterprise server 314 as necessary to retrieve information from the databases 316 pertinent to the

-240-

citation analysis tree function specified by the operator. The manner in which step 15712 is performed is illustrated in a flowchart shown in FIG. 158.

In step 15804, the network client 306 sends a request to the enterprise server 314. The network client 306 generates the request in accordance with the enterprise server API discussed above. Accordingly, the request generated by the network client 306 conforms to the enterprise server API. Preferably, the request is made using the ReqAnalysisCitation command (described above). Alternatively, the request can be made using a properly formed ReqFunction command (described above). The request specifies: (1) either a backward citation search or a forward citation search (as indicated by the user in dialogue box 16302); (2) the '484 patent; and (3) the number of levels (as specified by the operator in dialogue box 16302).

In step 15806, the enterprise server 314 processes the request. In particular, the enterprise server 314 accesses the patent table 1222 (FIG. 12H) to identify the document_id associated with the '484 patent. For backward citation searches, the enterprise server 314 then accesses the PatentRef Table 1208 (FIG. 12C) using this document_id as the key to identify all patents which are cited in the '484 patent. Each hit in the PatentRef table 1208 corresponds to a patent that is cited in the '484 patent. When performing a forward patent citation search, the enterprise server 314 accesses the PatentRef Table 1208 to identify all patents which cite the '484 patent. In this case, the enterprise server 314 uses the document_id of the '484 patent as a key in searching the RefPatentNo field of the PatentRef table 1208 (in backward searches, the document_id field of the PatentRef table 1208 is searched). Each hit in the PatentRef table 1208 corresponds to a patent that cites the '484 patent.

For patent citations of one level, the search is then complete. However, for patent citations of greater than one level, the enterprise server 314 must repeatedly accesses the Patent RefTable 1208. For example, consider a backward patent citation function of two levels. Suppose that the above described operation

-241-

identified that patents P1 and P2 are cited in the '484 patent. This is the first level. The second level refers to the patents that are cited in patents P1 and P2. Accordingly, the enterprise server 314 must perform the above described operations for each of patents P1 and P2, to identify the patents cited in P1 and P2.

In practice, to enhance performance, a single search command is issued for each level of the citation function. For example, a first search command is issued to identify all patents that cite a patent P1 (for a forward citation function). This is the first level. Suppose that patents P2, P3, and P4 are identified by this first search command. In the second level, a second search command is issued to identify all patents that cite patents P2, P3, and P4. In a similar manner, a single search command is issued for each of the other levels of the citation function being performed.

Further in step 15806, the enterprise server 314 processes the information retrieved from the databases 316 and generates two tables. In particular, the enterprise server 314 generates a parent/child table. An example patent/child Table 18002 is shown in FIG. 180. In the parent/child Table 18002, parent/child relationships between the patents retrieved from the databases 316 are indicated. The parent/child DAG 18102 corresponding to the parent/child Table 18002 is shown in FIG. 181. In DAG 18102, in the context of the example involving the '484 patent, node 1 corresponds to the '484 patent.

The enterprise server 314 also creates a patent bibliographic information table. An example patent bibliographic information Table 18202 is shown in FIG. 182. In the patent bibliographic information Table 18202, there is a row for each patent represented in the parent/child Table 18002. The row corresponding to each of these patents stores bibliographic information for that patent.

In other embodiments, the client 306 issues two commands, a first command to retrieve the parent/child information, and a second command to retrieve bibliographic information on the parents and children. In such

-242-

embodiments, the ReqAnalysisCitation function may be configured to optionally return only the parent/child information. Other enterprise server API commands (described above) could be used to retrieve patent bibliographic information.

5 In step 15808, the enterprise server 314 transfers the patent/child table 18002 and the patent bibliographic information table 18202 to the requesting network client.

10 Referring again to FIG. 157, in step 15714 the network client 306, after receiving the parent/child table 18002 and the patent bibliographic information table 18202, generates a citation analysis DAG corresponding to the patent/child relationships indicated in the parent/child table 18002. Alternatively, the patent bibliographic information table 18202 can be retrieved as necessary, when the operator wants to view a patent (as described below). As indicated above, FIG. 181 illustrates the citation analysis DAG 18102 for the parent/child table 18002 of FIG. 180. The manner in which the network client 306 performs step 15714 is represented by a flowchart shown in FIG. 159.

15 In step 15904, the network client 306 initializes the citation analysis DAG 18102. The citation analysis DAG 18102 includes a root node representing the selected patent (that is, the '484 patent), and nodes for each of the other patents (i.e., patents that cite the '484 patent, or that are cited in the '484 patent). In the example of FIG. 181, the root node, representing the '484 patent, is node 1. The other nodes 2-7 represent patents that are cited by the '484 patent (in the case of a backward citation function), or patents that cite the '484 patent (in the case of a forward citation function). Note that, upon completion of step 15904, the links between nodes in the citation analysis DAG 18102 have not yet been established.

20 In step 15906, the network client 306 selects one of the nodes in the citation analysis DAG 18102 for analysis.

25 In step 15908, the network client 306 generates a link for each parent/child relationship for the selected node, where the selected node is the child. Such parent/child relationship information is obtained from the parent/child Table

-243-

18002. For example, assume that node 7 is the selected node. As indicated in the parent/child Table 18002, node 7 (or, equivalently, patent 7) is the child of patent 5 and patent 6. Accordingly, in step 15908, the network client 306 creates a link to node 7 in both node 5 and node 6. The information representing these links is preferably stored in the parent nodes of the selected node to allow navigation of the tree from the root down.

In step 15910, the network client 306 determines whether there are any additional nodes in the citation analysis DAG 18102 that have not yet been processed. If there are such nodes, then control returns to step 15906. Otherwise, the citation analysis DAG 18102 is complete, and control returns to flowchart 15702 in FIG. 157.

Referring again to FIG. 157, in step 15716 the network client 306 creates the citation analysis hyperbolic tree 16404. When necessary (as described above), the network client 306 creates a tree (that differs from the hyperbolic tree 16404 that will ultimately be constructed) that corresponds to the citation analysis DAG 18102 generated in step 15714. The manner in which the network client 306 performs this function is represented by a flowchart in FIG. 160.

In step 16004, the network client 306 selects a node of the citation analysis DAG 18102 (it is noted that the root node 18104 corresponding to the '484 patent is not subject to selection). Also, the network client 306 initializes the tree 18302 (an example of which is shown in FIG. 183) by constructing the root node 18304 that corresponds to the '484 patent.

In step 16008, the network client 306 selects one of the parents of the selected node from the citation analysis DAG 18102. It is noted that, by definition, the selected node must have at least one parent. Assume, for example, that node 7 was selected in step 16004, and node 5 was selected in step 16008.

In step 16010, the network client 306 generates a new node in the tree 18302 being constructed. Specifically, the network client 306 generates node 7A in tree 18302.

-244-

In step 16012, the network client generates a link that represents the parent/child relationship between nodes 7A and 5. The information corresponding to this link is stored in tree node 5 of tree 18302.

5 In step 16014, the network client 306 determines whether the selected node (that is, node 7 in the citation analysis DAG 18102) has any additional parents in the DAG 18102 that have not yet been processed. In the case of the current example, node 7 has a parent, node 6, that has not yet been processed. Accordingly, control returns to step 16008 to process the parent node 6 with respect to the selected node 7. If, in step 16014, the selected node does not have
10 any additional parents in the DAG 18102 to process, then control flows to step 16016.

In step 16016, the network client 306 determines whether there are additional nodes in the citation analysis DAG 18102 to process. If there are additional nodes to process, then control returns to step 16004. Otherwise,
15 construction of tree 18302 is complete, and control returns to flowchart 15702 in FIG. 157.

Referring again to FIG. 157, further in step 15716 the tree 18302 is used to generate a hyperbolic tree. In generating the hyperbolic tree, a hyperbolic tree node is generated for each node in the tree 18302. Also, links between nodes in
20 the hyperbolic tree correspond to links in the tree 18302. In addition to linking information, the nodes in the hyperbolic tree may also include bibliographic information for corresponding patents. Such bibliographic information is obtained from the patent bibliographic information Table 18202. Alternatively, such patent bibliographic information is maintained in the patent bibliographic information
25 Table 18202, and retrieved as necessary (that is, such patent bibliographic information is not stored in the hyperbolic tree). Further details concerning construction of a hyperbolic tree corresponding to tree 18302 will be apparent to persons skilled in the relevant arts.

-245-

As discussed above, in some embodiments, the invention generates a hyperbolic tree directly from the DAG generated in step 15714. Whether this is possible depends on the capabilities of the hyperbolic browser being used in the client 306, and potentially other issues, such as performance issues. In these
5. embodiments of the invention where the hyperbolic tree is generated directly from the DAG, it is not necessary to perform the above described functions of step 15716.

In step 15718 the network client 306 displays the hyperbolic tree generated in step 15716. An example patent citation hyperbolic tree 16404 corresponding
10 to the '484 patent is shown in FIG. 164, wherein a portion of the patent citation tree 16404 is being displayed in a window 16412.

As indicated above, the nodes of the patent citation tree 16404 may be color coded or otherwise emphasized according to some criteria. Examples of such criteria are described below, although the invention is not limited to these
15 examples. The operation of step 15718 in performing this function is represented by a flowchart shown in FIG. 176.

In step 17604, based on user command, the network client 306 color codes or otherwise emphasizes and displays the nodes in the patent citation tree 16404 according to ownership and/or assignee (or some other criteria specified by the
20 operator). For example, the nodes corresponding to patents owned by company A may be color coded red, whereas the nodes corresponding to patents owned by company B may be color coded blue. Alternatively, the nodes corresponding to patents owned or licensed by the customer corporate entity may be color coded purple, whereas the nodes owned by any other corporate entity are color coded
25 black.

It is noted that other means for highlighting or emphasizing the nodes in the patent citation tree 16404 (other than color coding) can be used. In fact, any of the display properties of the nodes can be adjusted to reflect the criteria of

-246-

interest. Such display properties include uppercase, lowercase, typefaces, fonts, underlining, bold, italics, type size, etc.

5 The enterprise server 314 retrieves ownership and/or assignee information (upon receiving a command from the network client 306 requesting such information) from the assignee table 1201 (FIG. 12B) and/or the patent_ownership Table 1234 (FIG. 12J). This information is preferably returned to the network client 306 in the patent bibliographic information Table 18202 from the enterprise server 314.

10 In one embodiment, the assignment information in the assignee Table 1201 reflects any original assignment information on the front page of patents. In some embodiments, the assignee Table 1201 is periodically modified to reflect updated assignment information as recorded in the U.S. Patent Office or other national Patent Offices (in the case of non-U.S. patents).

15 In other embodiments, such updated assignment information is stored in a separate table (not shown), called the current_assignee table. The original assignment information in the assignee table 1201 is not changed.

20 In these other embodiments, according to the present invention, the operator can instruct the network client 306 to update the display of the patent citation tree 16404 to reflect such current assignment information using the information in the current_assignee table.

Accordingly, in step 17606, the network client 306 receives a command from the operator to update the patent citation tree 16404 with the current assignee information in the current_assignee table.

25 In step 17608, the network client 306 sends a request to the enterprise server 314 for updated assignment information. This request, which is in conformance with the enterprise server API, lists the patents represented in the patent citation tree 16404. In accordance with this request, the enterprise server 314 accesses the current_assignee table for any updated assignment information regarding the listed patents. It is noted that not all patents may have updated

-247-

assignment information. The enterprise server 314 transfers this information back to the requesting network client 306.

5 In step 17610, the network client 306 adjusts the display properties (such as the color coding) of the nodes in the patent citation tree 16404 to reflect any new assignment information as received from the enterprise server 314. According to an embodiment of the invention, the network client 306 allows the operator to switch back and forth between the original assignment information and the updated assignment information. In other words, upon appropriate operator command, the network client 306 toggles between displaying the patent citation
10 tree 16404 according to the assignment information stored in the assignee Table 1201, and displaying the patent citation tree 16404 according to the updated assignee information in the current_assignee table.

Nodes can also be color coded according to other criteria as selected by the operator. For example, in step 17604, based on user command, the network
15 client 306 color codes or otherwise emphasizes and displays the nodes in the patent citation tree 16404 according to issue date. For example, patents that issued in 1984-1986 are color coded blue, patents that issued in 1986-1988 are color coded red, patents that issued in 1988-1990 are color coded yellow, etc. In a similar manner, the nodes can be color coded according to other time-based
20 criteria, such as filing date, priority date, length of pendency, effective filing date, invention date, critical date, on-sale date, public disclosure date, public use date, etc. Such information is either retrieved from the databases 316, calculated based on information retrieved from the databases 316, or entered by operators. Information entered by an operator on a patent (such as the invention date) is
25 preferably stored in the node that is associated with the patent. Such information is, upon receipt of appropriate user command, referred to when displaying and emphasizing the node.

Operators can also specify the time increment (1984-1986 or 1984-1988, for example), the beginning date to view, the ending date to view, etc.

-248-

The invention also allows operators to enter their own user-defined criteria. For example, an operator can specify that patents P1 and P3 should be color coded blue, patents P10 and P21 should be color coded red, etc. Such specification by the operator may be based on information or analysis conducted by the operator. For example, the operator may have previously determined that patents P1 and P3 are unenforceable, and patents P10 and P21 contain claims that are infringed.

In an embodiment, the operator right clicks over a node to display a context sensitive menu. Via this menu, the operator can specify user-defined criteria for displaying the node (i.e., that the node should be color coded blue, or the node should be color coded as an unenforceable patent, where the color for unenforceable patents is defined elsewhere). The operator has the option of entering text that describes the meaning of his criteria, such as "Blue: Patents that are infringed." This text is displayed in the legend (described below).

Preferably, nodes can be emphasized using any combination of operator-selected criterions. This is done by using different display properties to represent different criterions. For example, color could be used to represent ownership, font could be used to represent issue date, underlining could be used to represent enforceable/unenforceable, etc.

According to an embodiment of the invention, a legend is displayed next to the patent citation tree 16404. This legend specifies the meaning for the colors or other display properties used to highlight the nodes of the patent citation tree 16404. For example, the legend may indicate that blue is for patents owned by IBM, or is for patents issued in 1994-1997, or is for patents indicated by the operator as being infringed.

The present invention enables operators to interact with the data in the patent citation tree 16404. For example, operators can elect to view information regarding any of the patents represented in the patent citation tree 16404. Such operation is represented by steps 15720, 15722, and 15724.

-249-

In step 15720, the operator selects one of the patents represented in the patent citation tree 16404. In the example of FIG. 164, the operator has selected the '168 patent by double-clicking while the mouse (represented by icon 16412) is positioned over node 16408 corresponding to the '168 patent.

5 In step 15722, the network client 306 retrieves information pertaining to the selected '168 patent.

In step 15724, at a minimum, the network client 306 displays the Abstract and bibliographic information on the '168 patent. This information was preferably transferred to the network client 306 in the patent bibliographic information table 18202 (FIG. 182). FIG. 165 illustrates an example window 16502 where the Abstract and bibliographic information pertaining to the 186 patent is displayed.

10 If the selected patent (that is, the '168 patent) is in the local repository, then the network client 306 in step 15722 interacts with the enterprise server 314 to download the patent text and/or the patent image of the '168 patent. The network client 306 then displays the patent text and/or patent image as described above (see, for example, FIGS. 123 and 124). This is represented by step 15724.

15

start here

As indicated above, the operator can elect to color code or otherwise adjust the display properties of the nodes in the patent citation tree 16404 according to various criteria. Such criteria includes assignee and ownership information, as described above. Another criteria is patent aging. Specifically, an operator can elect to adjust the display properties of the nodes in the patent citation tree 16404 according to the age of the patents relative to the base patent.

20 In the example of FIG. 164, the base patent is the '484 patent which was originally selected by the operator in step 15706 (FIG. 157). According to this aspect of the invention, the nodes corresponding to the other patents in the patent citation tree 16404 are color coded to reflect their age relative to the base '484 patent. For example, patents that issued less than a year after the '484 patent are color coded blue. Patents that issued more than one year but less than two years

25

-250-

after the '484 patent are color coded red. Patents that issued more than two years but less than three years after the '484 patent are color coded yellow.

Alternatively, gray scales can be used to reflect the age of the patents relative to the base '484 patent. For example, the closer the issue date of a patent is to the issue date of the '484 patent, the darker the gray used, and the farther away the issue date of a patent is to the issue date of the '484 patent, the lighter the gray used.

Alternatively, patent aging can be visualized by displaying in the patent citation tree the root plus any patent that issued within a year of the root (or any other time increment specified by the operator). The tree then grows upon receipt of user command to include patents that issued later. For example, upon receipt of user command, the tree grows to additionally include patents that issued within two years of the root, then three years, four years, five years, etc.

Patent aging visualizations (described here and in other places of this application) are valuable because they represent when patent-related activity occurred with respect to a given technology.

Patent Citation Tree (Web Client)

The operation of the patent citation tree function when performed by a web client 304 interacting with the enterprise server 314 via the web server 310 (FIG. 3) shall now be described with reference to a flowchart 16602 shown in FIG. 166.

In step 16606, the operator at the web client 304 defines a search to be performed by entering search criteria into a patent search window, such as patent search widow 14002 shown in FIG. 140. In the manner described above, the web client 304 interacts with the enterprise server 314 via the web server 310 to conduct the search defined by the operator. The web client 304 displays the

-251-

search results in a search results window, such as search results widow 14102 shown in FIG. 141. This is represented by step 16608 in flowchart 16602.

In step 16610, the operator selects one of the patents listed in the search results window 14102. In step 16612, the web client 304 interacts with the enterprise server 314 via the web server 310 to retrieve information pertaining to the selected patent, and to display such information (such processing is described above). Preferably, such information is displayed in windows, such as those shown in FIGS. 144, 145A-145C, and/or FIG. 146. Further details regarding the operation of steps 16606, 16608, 16610, and 16612 are provided above.

In step 16614, the operator at the web client 304 issues a command to perform a patent citation analysis of the patent being displayed. For example, suppose that in step 16612, information regarding an '830 patent was retrieved and displayed in a window 14602 shown in FIG. 146. In step 16614, the operator presses a patent citation tree button 14610 in order to perform a patent citation analysis of the '830 patent.

In step 16616, the web client 304 presents a dialog box 16302 (FIG. 163) to the operator. In this dialog box 16302, the operator identifies whether he wishes to perform a forward citation function or a backward citation function. Also, the operator defines the scope of the citation function to be performed by indicating the number of levels that should be cited. This information entered into the dialog box 16302 is received by the web client 304.

In step 16618, the web client 304 interacts with the enterprise server 314 via the web server 310 to retrieve information from the databases 316 pertinent to the citation analysis tree function specified by the operator. The manner in which step 16618 is performed is illustrated in a flowchart shown in FIG. 167.

In step 16704, the browser 808 in the web client 304 generates and sends a request to the web server 310 (see FIG. 8). This request preferably includes patent-centric URL commands as described above. The request specifies: (1) either a backward citation search or a forward citation search (as indicated by

-252-

the operator in dialog box 16302); (2) the '830 patent; and (3) the number of levels (as specified by the operator in dialog box 16302). Preferably, the request is formed using the GetCitationTree command, described above.

5 In step 16706, the web server 310 receives the request from the browser 808. The translator 804 in the web server 310 translates the request to a new request that is in conformance with the enterprise server API. This translation function is described above. The web server 310 sends the translated request to the enterprise server 314. In effect, the web server 310 is playing the role of the network client 306.

10 In step 16708, the enterprise server 314 processes the request. In particular, the enterprise server 314 access the patent table 12222 (FIG. 12H) to identify the document_id associated with the '830 patent. For backward citation searches, the enterprise server 314 then accesses the PatentRef table 1208 to identify all patents which are cited in the '830 patent. For forward patent citation
15 searches, the enterprise server 314 accesses the PatentRef table 1208 to identify all patents which cite the '830 patent. The enterprise server '314 repeatedly accesses the PatentRef table 1208 according to the number of levels as indicated in the request from the web server 310. Further details regarding the operation of the enterprise server 314 are described above.

20 The enterprise server 314 processes the information retrieved from the databases 316 and generates two tables, the parent child table 18002 (FIG. 180) and the patent bibliographic information table 18202 (FIG. 182). The operation of the enterprise server 314 when performing step 16708 is similar to its operation when performing step 15806 (described above).

25 In step 16710, the enterprise server 314 transmits the parent/child table 18002 and the patent bibliographic information table 18202 to the web server 310.

In step 16712, the translator 804 in the web server 310 translates, as necessary, the information received from the enterprise server 314, and forwards the translated information to the web client 304. The translator 804 translates the

-253-

information from the enterprise server 314 to a format that is compatible with the capabilities and plug-ins associated with the browser 808. This operation shall be further described with reference to FIG. 186.

As shown in FIG. 186, in an embodiment of the invention, the web client
5 304 includes a hyperbolic tree module 18602 that interacts with the browser 808. In some embodiments, the hyperbolic tree module 18602 may optionally include or interact with a hyperbolic tree DLL (dynamic link library), not shown.

The translator 804 in step 16712 translates the information received from
10 the enterprise server 314 to a format that is useable by the hyperbolic tree module 18602. Accordingly, the operation of the translator 804 in this respect is specific to the particular functions and capabilities of the hyperbolic tree module 18602 being used in the web client 304. The browser 808 receives this translated information from the web server 310, and then forwards the translated information to the hyperbolic tree module 18602 for processing. As described below, the
15 hyperbolic tree module 18602 modifies and formats this information from the browser 808, as necessary, and then displays and processes the patent citation hyperbolic tree.

Referring again to FIG. 166, in step 16620, the hyperbolic tree module
20 18602 in the web client 304 generates a citation analysis DAG from the information received from the enterprise server 314 via the web server 310. The operation of the hyperbolic tree module 18602 when performing step 16620 is similar to the operation of the network client 306 when performing step 15714.

In step 16622, the hyperbolic tree module 18602 maps the citation analysis
25 DAG to a tree where necessary (as described above), and then generates a hyperbolic tree from this tree. The operation of the hyperbolic tree module 18602 when performing step 16622 is similar to the operation of the network client 306 when performing step 15716.

In some embodiments, at least parts of steps 16620 and 16622 are performed by the translator 804 in the web server 310. In particular, the translator

-254-

804 generates a citation analysis DAG from the information received from the enterprise server 314, and maps the citation analysis DAG to a tree. The translator 804 may also generate a hyperbolic tree from the tree. The translator 804 forwards the results of its processing (i.e., the translated information referred to above) to the browser 808.

As discussed above, in some embodiments, the invention generates a hyperbolic tree directly from the DAG generated in step 16620. Whether this is possible depends on the capabilities of the hyperbolic tree module 18602 in the web client 304, and potentially other issues, such as performance issues. In these embodiments of the invention where the hyperbolic tree is generated directly from the DAG, it is not necessary to perform the above described functions of step 16622.

In step 16624, the hyperbolic tree module 18602 displays the patent citation information to the operator in a hyperbolic tree format. In other words, in step 16624, the hyperbolic tree module 18602 displays the hyperbolic tree generated in step 16622. The operation of the hyperbolic tree module 18602 when performing step 16624 is similar to the operation of the network client 306 when performing step 15718, wherein at the web client 304, the patent citation tree 16404 is displayed in a window similar to the window 16412 shown in FIG. 164.

In step 16626, the operator at the web client 304 selects one of the patents represented in the patent citation tree 16404. Step 16626 is similar to step 15720 (FIG. 157).

In step 16628, the web client 304 retrieves information on the patent selected in step 16626. Step 16628 is similar to step 15722 (FIG. 157).

In step 16630, the web client 304 displays the retrieved patent information. Step 16630 is similar to step 15724 (FIG. 157).

Preferably, the web client 304, when displaying the patent citation tree 16404, color codes or otherwise emphasizes the nodes of the patent citation tree

16404 according to some user selected criteria, such as ownership, assignee, relative age, issue date, filing date, user-defined criteria, etc., as described above.

Additional Patent Citation Visualizations

5 The present invention includes the capability to display patent citation data in additional formats. For example, FIG. 174 illustrates an example display format 17402 where patent citation information is visualized as a modified bar graph. Display format 17402 represents a backward patent citation visualization of a '201 patent. The '201 patent is cited by a '567 patent, a '432 patent, a '123 patent, and an '888 patent. Each of these patents is represented by a bar. The horizontal
10 length of each bar graphically represents the pendency duration of the associated patent. The longer the bar, the longer the pendency of the patent. Also, the beginning and ending locations of each bar graphically represent the filing and issue dates of the patent. For example, the '201 patent was filed at the beginning of 1989, and issued at the end of 1989. The bars in display format 17402 can be
15 color-coded or otherwise emphasized to reflect user-selected criteria, such as ownership, assignee, etc., as described above.

Note that display format 17402 in FIG. 174 represents a one-level patent citation function. Additional levels can be viewed by selecting the patent of interest. For example, if the operator double clicks on the '123 patent, then the
20 patents that cite the '123 patent will be displayed. This is shown in FIG. 175.

Patent Claims Tree

According to an embodiment of the present invention, the analysis modules 416 in the enterprise server 314 include a patent claim tree module 1022 (FIG. 10). FIG. 171 illustrates an example patent claims tree 17101. In this
25 patent claims tree 17101, the root node 17104 represents the patent (in this case,

-256-

U.S. Patent No. 5,092,011). The nodes 17106 directly connected to the root node 17104 represent the independent claims in the patent. The nodes connected to these independent claim nodes represent the dependent claims of the respective independent claims. The links in the patent claims tree 17101 represent the dependencies between claims.

The nodes in the patent claims tree 17101 can be color coded or otherwise emphasized (by modifying some display property) according to various criteria, including user-defined criteria. For example, when performing an infringement analysis, the operator can elect to color code claims that are believed to be infringed using one color, and claims that are not believed to be infringed using another color. When conducting a validity analysis, the operator can elect to color code claims that are believed to be valid using one color, and claims that are believed to be invalid using a different color. As discussed below, nodes can also be color coded or emphasized according to whether they correspond to patents or claims, or whether they correspond to independent or dependent claims. Also, as with the patent citation tree (described above), operators can color code according to operator-defined criteria. Other criterions for color coding the claims will be apparent to persons skilled in the relevant art(s).

The patent claims tree function, when performed by either a network client 306 or a web client 304, shall now be described with reference to a flowchart 16802 shown in FIG. 168.

In step 16806, the operator selects a patent using any of the mechanisms described herein, such as selecting a patent from a listing of patents displayed in the document pane 16104 of the console 16102 (FIG. 161).

In step 16808, information pertaining to the selected patent is retrieved from the databases 316 (as necessary) and displayed using any of the display formats described herein (see, for example, FIGS. 123 and 124).

In step 16810, the operator issues a command to display a patent claims tree for the patent being displayed. The operator may issue this command using

-257-

any of the mechanisms described herein, such as by selecting a command from the menu bar, selecting a command from a pop-up window, selecting the command from a toolbar, etc.

5 In step 16812, the client 304/306 retrieves (as necessary) and analyzes the text of the claims section of the patent in order to identify claim dependencies (since it is necessary to analyze the claims, the patent claims tree function according to one embodiment of the invention is only available for those patents in the local repository). The dependency relationship between claims is identified by parsing the claim text of the patent and identifying claim dependency language
10 contained therein (particularly in the claim preambles), such as "the apparatus of claim 1," "the method of claim 1," "the method according to claims 3 or 4," "the system as in one of claims 4-7," "a product produced according to the method of claim 1," "a system, comprising," etc. The dependency between claims is derived by inspection and analysis of these phrases in the preambles of the claims.

15 In step 16814, the client 304/306 generates a claims dependency DAG. For example, FIG. 184 illustrates a claims dependency DAG 18402 corresponding to the '011 patent shown in FIG. 171. The manner in which the client 304/306 performs step 16814 is represented in a flowchart shown in FIG. 169.

20 In step 16904, the client 304/306 initializes the claims dependency DAG 18402. This initial DAG 18402 includes a root node representing the '011 patent, and a node for each claim in the '011 patent.

In step 16906, the client 304/306 selects one of the nodes of the DAG 18402 (or, equivalently, the client 304/306 selects one of the claims of the '011 patent).

25 In step 16908, if the selected claim is an independent claim, then the client 304/306 creates a link from the root node to the selected node. For example, if claim 1 is the selected claim, then the client 304/306 creates a link between node 18406 to the root node 18408. Information representative of the link is stored in

-258-

the parent of the selected claim. For example, information representing the link between node 18406 and the root node 18408 is stored in root node 18408.

If the selected claim is a dependent claim, then in step 16908 the client 304/306 creates a link between the selected node and each node (called parent nodes) corresponding to a claim from which the selected claim depends. Information representing these links is stored in the parent nodes. For example, assume that node 18404 corresponding to claim 10 was selected in step 16906. Claim 10 depends from claims 1 and 9, the parent nodes of claim 10. Accordingly, in step 16908, the client 304/306 creates a link from node 18404 to node 18406 and node 18410. The link information is stored in parent nodes 18406 and 18410.

In step 16910, the client 304/306 determines whether there are additional claims in the '011 patent to process. If there are additional claims to process, then control returns to step 16906. Otherwise, construction of the claims dependency DAG 18402 is complete, and control returns to flowchart 16802.

Referring again to FIG. 168, in step 16816 the client 304/306 generates the patent claims hyperbolic tree. In so doing, the client 304/306 generates a claims dependency tree, if necessary (as described above). FIG. 185 illustrates a claims dependency tree 18502 that corresponds to the claims dependency DAG 18402 in FIG. 184. The manner in which client 304/306 generates the claims dependency tree 18502 is represented by a flowchart in FIG. 170.

In step 17004, the claims dependency tree 18502 is initialized to include the root node 18504 corresponding to patent '011. Also, the client 304/306 selects a node from the claims dependency DAG 18402 for processing. Preferably, the root node 18408 is not subject to selection (since a corresponding root node 18504 in the claims dependency tree 18502 has already been created). Preferably, nodes in the claims dependency DAG 18402 are selected in ascending order of their corresponding claims. For example, the node corresponding to claim 1 is selected first. The nodes corresponding to claims 2, 3, 4, etc., are then selected. In this manner, it is assured that, in the claims dependency tree 18502,

-259-

the tree nodes corresponding to a dependent claim's parent claims will exist before the dependent claim is processed.

In step 17008, the client 304/306 selects one of the parents of the selected node. By definition, each node (other than the root node) in the claims dependency DAG 18402 includes at least one parent.

In step 17010, the client 304/306 creates a new node in the claims dependency tree 18502. This new tree node corresponds to the selected node of the claims dependency DAG 18402. Any claim information (such as the text of the claim) is optionally stored in this new tree node.

In step 17012, in the claims dependency tree 18502, the client 304/306 generates a link between the node corresponding to the selected parent, and the new tree node. Information for this link is stored in the selected parent node. For example, assume in step 17004 that node 18404 corresponding to claim 10 was selected, and in step 17008 claim 1 (parent to claim 10) was selected. In step 17012, a link is generated between node 18404A (corresponding to claim 10) and node 18510 (corresponding to claim 1) in the claims dependency tree 18502. The link information is preferably stored in node 18510.

In step 17014, the client 304/306 determines whether the selected node has any additional parents in the claims dependency DAG 18402 that has not yet been processed. If the selected node in the claims dependency DAG 18402 has additional parents, then control returns to step 17008 to process another parent. Otherwise, step 17016 is performed.

In step 17016, the client 304/306 determines whether there are additional nodes in the claims dependency DAG 18402 to process. If there are additional nodes to process, then control returns to step 17004. Otherwise, construction of the claims dependency tree 18502 is complete.

Referring again to FIG. 168, further in step 16816, the client 304/306 uses the claims dependency tree 18502 to generate the patent claims hyperbolic tree 17101. In generating the hyperbolic tree, a hyperbolic tree node is generated for

-260-

each node in the claims dependency tree 18502. Also, links between nodes in the patent claims hyperbolic tree 17101 correspond to links in the claims dependency tree 18502. In addition to linking information, the nodes in the patent claims hyperbolic tree 17101 may also include any claim information that was in the corresponding node of the claims dependency tree 18502, such as claim text. The node type and the sub-type are also stored in the nodes of the patent claims hyperbolic tree 17101. The node type refers to whether or not the node is a patent node (such as node 17104) or a claim node (such as node 17108A). The sub-type refers to whether the node corresponds to an independent claim (such as nodes 17106) or a dependent claim (such as nodes 17108A and 17108B). Further details concerning construction of a hyperbolic tree representative of the claims dependency tree 18502 will be apparent to persons skilled in the relevant art(s).

As discussed above, in some embodiments, the invention generates a hyperbolic tree directly from the DAG generated in step 16814. Whether this is possible depends on the capabilities of the hyperbolic browser being used in the client 304/306, and potentially other issues, such as performance issues. In these embodiments of the invention where the hyperbolic tree is generated directly from the DAG, it is not necessary to perform the above described functions of step 16816.

In step 16818, the client 304/306 displays the patent claims tree 17101 generated in step 16816. As indicated above, an example patent claims tree 17101 is shown in FIG. 171.

In displaying the patent claims hyperbolic tree 17101 in step 16818, the nodes of the patent claims tree 17101 may be color coded or otherwise emphasized (by modifying display properties) according to a number of criteria, such as node type and/or sub-type. For example, patent nodes may be color coded a first color whereas claim nodes may be color coded a second color. Nodes corresponding to independent claims may be color coded a third color, and nodes corresponding to dependent claims may be color coded a fourth color. As described

-261-

above, nodes can also be color coded according to user-defined criteria and user-provided information, such as claims that are being infringed, claims that are invalid, etc.

5 The present invention enables operators to interact with the data displayed in the patent claims tree 17101. For example, operators can elect to view the text of claims of any of the claims represented in the patent claims tree 17101. Such operation is represented by steps 16820 and 16822.

In step 16820, the operator selects one of the patents in the patent claims tree 17101 using any of the selection mechanisms described above.

10 In step 16822, the client 304/306 retrieves information on the selected claim, such as the text of the selected claim, and displays the claim text. The invention includes a number of formats for displaying the claim text. For example, FIG. 172 illustrates a display format 17202 where the text of the claim is displayed independent of and apart from the patent. In contrast, FIG. 173 illustrates a display format 17302 where the claim is displayed in the context of the patent.

15 Such user interaction may be based on the node type and subtype. For example, if the user selects a patent node, such as node 17104, then the patent would be displayed (see FIGS. 123 and 124). Alternatively, if the operator selects a node corresponding to a claim, then the text of the claim is displayed, as
20 described above.

Conclusion

25 While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What Is Claimed Is:

1. A method of processing data, comprising the steps of:
 - (1) maintaining at least one first database of patents;
 - (2) maintaining at least one second database of non-patent information
5 of interest to a corporate entity;
 - (3) maintaining one or more groups, each of said one or more groups comprising any number of said patents from said at least one first database; and
 - (4) automatically processing said patents in one of said one or more groups in conjunction with non-patent information from said at least one second
10 database.
2. The method of claim 1, further comprising the following steps that are performed before step (3):
 - (a) associating one of said one or more groups with a product;
 - (b) identifying patents from said at least one first database that map to
15 said product; and
 - (c) adding said identified patents to said one of said one or more groups.
3. The method of claim 2, further comprising the steps of:
 - (d) determining relevance rankings for said identified patents based on
20 a relation of said product to said corporate entity; and
 - (e) storing said relevance rankings in said at least one second database.
4. The method of claim 2, wherein said product is one that is being currently produced by said corporate entity, or is one that is being considered for future production by said corporate entity.

-263-

5. The method of claim 2, wherein step (a) comprises the step of:
associating said one of said one or more groups with a BOM (bill of materials) node in a BOM structure, wherein said BOM node represents an abstraction of said product.

5 6. The method of claim 2, wherein step (b) comprises the step of:
identifying patents from said at least one first database that map to said product, wherein said identified patents each includes at least one claim that potentially reads on said product.

10 7. The method of claim 1, wherein step (4) comprises the step of:
(a) automatically performing a patent mapping function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

(i) receiving from an operator information that identifies one of said one or more groups;

15 (ii) identifying patents that are within said identified group;

(iii) retrieving information pertaining to said identified patents from at least one of (I) said at least one first database and (II) said at least one second database; and

(iv) displaying said retrieved information.

20 8. The method of claim 7, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

9. The method of claim 1, wherein step (4) comprises the step of:
(a) automatically performing a backwards patent citation function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

25

-264-

(i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying first patents that are within said identified group;

(iii) identifying second patents that are cited in said first patents;

5 (iv) retrieving information pertaining to said second patents from at least one of (I) said at least one first database and (II) said at least one second database; and

(v) displaying said retrieved information.

10 10. The method of claim 9, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

11. The method of claim 1, wherein step (4) comprises the step of:

(a) automatically performing a forwards patent citation function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

15 (i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying first patents that are within said identified group;

(iii) identifying second patents that cite said first patents; ~~which are~~

20 (iv) retrieving information pertaining to said second patents from at least one of (I) said at least one first database and (II) said at least one second database; and

(v) displaying said retrieved information.

12. The method of claim 11, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

25 13. The method of claim 1, wherein step (4) comprises the step of:

-265-

(a) automatically performing a patent aging function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

5 (i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying patents that are within said identified group;

(iii) retrieving patent term expiration related information for said identified patents from said at least one first database; and

10 (iv) displaying patent aging information for said identified patents based on said retrieved patent term expiration related information.

14. The method of claim 13, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

15. The method of claim 1, wherein step (4) comprises the step of:

15 (a) automatically performing a backwards patent bracketing/clustering function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

(i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying first patents that are within said identified group;

20 (iii) identifying second patents that are cited in said first patents;

(iv) retrieving ownership information pertaining to said first and second patents from said at least one second database; and

25 (v) displaying a connected graph of said first patents and said second patents, wherein said connected graph depicts said retrieved ownership information.

-266-

16. The method of claim 15, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

17. The method of claim 1, wherein step (4) comprises the step of:

(a) automatically performing a forwards patent bracketing/clustering function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

(i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying first patents that are within said identified group;

(iii) identifying second patents that cite said first patents;

(iv) retrieving ownership information pertaining to said first and second patents from said at least one second database; and

(v) displaying a connected graph of said first patents and said second patents, wherein said connected graph depicts said retrieved ownership information.

18. The method of claim 17, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

19. The method of claim 1, wherein step (4) comprises the step of:

(a) automatically performing an inventor patent count function with respect to said patents in one of said one or more groups, wherein step (a) comprises the steps of:

(i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying patents that are within said identified group;

(iii) retrieving inventorship information pertaining to said identified patents from said at least one second database;

-267-

(iv) determining top inventors of said identified patents based on said retrieved inventorship information; and

(v) displaying information pertinent to said top inventors.

20. The method of claim 19, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

21. The method of claim 1, wherein step (4) comprises the step of:

(a) automatically performing an inventor employment information function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

(i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying patents that are within said identified group;

(iii) identifying inventors of said identified patents by retrieving information pertaining to said identified patents from said at least one second database;

(iv) retrieving inventorship information pertaining to said identified inventors from said at least one second database; and

(v) displaying said retrieved inventorship information.

22. The method of claim 21, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

23. The method of claim 1, wherein step (4) comprises the step of:

(a) automatically performing a financial related function with respect to patents in one of said one or more groups, wherein step (a) comprises the steps of:

-268-

(i) receiving from an operator information that identifies one of said one or more groups;

(ii) identifying patents that are within said identified group;

5 (iii) retrieving financial information pertaining to said identified patents from said at least one second database; and

(iv) displaying said retrieved financial information.

24. The method of claim 1, further comprising the following steps that are performed before step (3):

(a) associating one of said one or more groups with a person;

10 (b) identifying patents from said at least one first database that name as an inventor said person; and

(c) adding said identified patents to said one of said one or more groups.

15 25. The method of claim 1, further comprising the following steps that are performed before step (3):

(a) associating one of said one or more groups with a second corporate entity;

(b) identifying patents from said at least one first database that are associated with said second corporate entity; and

20 (c) adding said identified patents to said one of said one or more groups.

26. The method of claim 25, wherein step (b) comprises the step of:
identifying patents from said at least one first database that are owned or licensed by said second corporate entity.

25 27. A system of processing data, comprising:

-269-

at least one first database of patents;

at least one second database of non-patent information of interest to a corporate entity;

group maintaining means for maintaining one or more groups, each of said one or more groups comprising any number of said patents from said at least one first database; and

processing means for automatically processing said patents in one of said one or more groups in conjunction with non-patent information from said at least one second database.

28. The system of claim 26, wherein one of said one or more groups are associated with a product, further comprising:

patent identification assisting means for assisting in identifying patents from said at least one first database that map to said product; and

means for adding said identified patents to said one of said one or more groups.

29. The system of claim 28, further comprising:
means for assisting in determining relevance rankings for said identified patents based on a relation of said product to said corporate entity; and

means for storing said relevance rankings in said at least one second database.

30. The system of claim 28, wherein said product is one that is being currently produced by said corporate entity, or is one that is being considered for future production by said corporate entity.

-270-

31. The system of claim 28, wherein said one of said one or more groups is associated with a BOM (bill of materials) node in a BOM structure, wherein said BOM node represents an abstraction of said product.

5 32. The system of claim 28, wherein said patent identification assisting means comprises:

means for assisting in identifying patents from said at least one first database that map to said product, wherein said identified patents each includes at least one claim that potentially reads on said product.

10 33. The system of claim 27, wherein said processing means comprises:

a patent mapping module to automatically perform a patent mapping function with respect to patents in one of said one or more groups, said patent mapping module comprising:

means for receiving from an operator information that identifies one of said one or more groups;

15 means for identifying patents that are within said identified group;

means for retrieving information pertaining to said identified patents from at least one of (I) said at least one first database and (II) said at least one second database; and

means for displaying said retrieved information.

20 34. The system of claim 33, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

35. The system of claim 27, wherein said processing means comprises:

a patent citation module to automatically perform a backwards patent citation function with respect to patents in one of said one or more groups, said patent citation module comprising:

25

-271-

means for receiving from an operator information that identifies one of said one or more groups;

means for identifying first patents that are within said identified group;

5 means for identifying second patents that are cited in said first patents;

means for retrieving information pertaining to said second patents from at least one of (I) said at least one first database and (II) said at least one second database; and

10 means for displaying said retrieved information.

36. The system of claim 35, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

37. The system of claim 27, wherein said processing means comprises:

15 a patent citation module to automatically perform a forwards patent citation function with respect to patents in one of said one or more groups, said patent citation module comprising:

means for receiving from an operator information that identifies one of said one or more groups;

20 means for identifying first patents that are within said identified group;

means for identifying second patents that cite said first patents;

means for retrieving information pertaining to said second patents from at least one of (I) said at least one first database and (II) said at least one second database; and

25 means for displaying said retrieved information.

-272-

38. The system of claim 37, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

39. The system of claim 27, wherein said processing means comprises:
a patent aging module to automatically perform a patent aging function
5 with respect to patents in one of said one or more groups, said patent aging module comprising:

means for receiving from an operator information that identifies one of said one or more groups;

means for identifying patents that are within said identified group;

10 means for retrieving patent term expiration related information for said identified patents from said at least one first database; and

means for displaying patent aging information for said identified patents based on said retrieved patent term expiration related information.

40. The system of claim 39, wherein said identified group is one of a BOM
15 group, a corporate entity group, an inventor group, and a user-defined group.

41. The system of claim 27, wherein said processing means comprises:
a patent bracketing/clustering module to automatically perform a
backwards patent bracketing/clustering function with respect to patents in one of
said one or more groups, said patent bracketing/clustering module comprising:

20 means for receiving from an operator information that identifies one of said one or more groups;

means for identifying first patents that are within said identified group;

25 means for identifying second patents that are cited in said first patents;

-273-

means for retrieving ownership information pertaining to said first and second patents from said at least one second database; and

means for displaying a connected graph of said first patents and said second patents, wherein said connected graph depicts said retrieved ownership information.

42. The system of claim 41, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

43. The system of claim 27, wherein said processing means comprises:
a patent bracketing/clustering module to automatically perform a forwards patent bracketing/clustering function with respect to patents in one of said one or more groups, said patent bracketing/clustering module comprising:

means for receiving from an operator information that identifies one of said one or more groups;

means for identifying first patents that are within said identified group;

means for identifying second patents that cite said first patents;
means for retrieving ownership information pertaining to said first and second patents from said at least one second database; and

means for displaying a connected graph of said first patents and said second patents, wherein said connected graph depicts said retrieved ownership information.

44. The system of claim 43, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

45. The system of claim 27, wherein said processing means comprises:

-274-

an inventor patent count module to automatically perform an inventor patent count function with respect to said patents in one of said one or more groups, said inventor patent count module comprising:

- 5 means for receiving from an operator information that identifies one of said one or more groups;
- means for identifying patents that are within said identified group;
- means for retrieving inventorship information pertaining to said identified patents from said at least one second database;
- 10 means for determining top inventors of said identified patents based on said retrieved inventorship information; and
- means for displaying information pertinent to said top inventors.

46. The system of claim 45, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

47. The system of claim 27, wherein said processing means comprises:

- 15 an inventor employment information module to automatically perform an inventor employment information function with respect to patents in one of said one or more groups, said inventor employment information module comprising:
- means for receiving from an operator information that identifies one of said one or more groups;
- 20 means for identifying patents that are within said identified group;
- means for identifying inventors of said identified patents by retrieving information pertaining to said identified patents from said at least one second database;
- 25 means for retrieving inventorship information pertaining to said identified inventors from said at least one second database; and
- means for displaying said retrieved inventorship information.

-275-

48. The system of claim 47, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

49. The system of claim 27, wherein said processing means comprises:
a financial module to automatically perform a financial related function
5 with respect to patents in one of said one or more groups, said financial module comprising:

means for receiving from an operator information that identifies
one of said one or more groups;

means for identifying patents that are within said identified group;

10 means for retrieving financial information pertaining to said
identified patents from said at least one second database; and

means for displaying said retrieved financial information.

50. The system of claim 27, wherein one of said one or more groups is
associated with a person, further comprising:

15 means for assisting in identifying patents from said at least one first
database that name as an inventor said person; and

means for adding said identified patents to said one of said one or more
groups.

51. The system of claim 27, wherein one of said one or more groups is
20 associated with a second corporate entity, further comprising:

means for assisting in identifying patents from said at least one first
database that are associated with said second corporate entity; and

means for adding said identified patents to said one of said one or more
groups.

-276-

52. The system of claim 51, wherein said identified patents are owned or licensed by said second corporate entity.

53. A computer program product comprising a computer useable medium having computer program logic stored therein, said computer program logic comprising:

means for enabling a computer to access at least one first database of patents;

means for enabling the computer to access at least one second database of non-patent information of interest to a corporate entity;

group maintaining means for enabling the computer to maintain one or more groups, each of said one or more groups comprising any number of said patents from said at least one first database; and

processing means for enabling the computer to automatically process said patents in one of said one or more groups in conjunction with non-patent information from said at least one second database.

54. The computer program product of claim 53, wherein one of said one or more groups are associated with a product, said computer program logic further comprising:

patent identification assisting means for enabling the computer to assist in identifying patents from said at least one first database that map to said product; and

means for enabling the computer to add said identified patents to said one of said one or more groups.

55. The computer program product of claim 54, said computer program logic further comprising:

-277-

means for enabling the computer to assist in determining relevance rankings for said identified patents based on a relation of said product to said corporate entity; and

means for enabling the computer to store said relevance rankings in said at least one second database.

56. The computer program product of claim 54, wherein said product is one that is being currently produced by said corporate entity, or is one that is being considered for future production by said corporate entity.

57. The computer program product of claim 54, wherein said one of said one or more groups is associated with a BOM (bill of materials) node in a BOM structure, wherein said BOM node represents an abstraction of said product.

58. The computer program product of claim 54, wherein said patent identification assisting means comprises:

means for enabling the computer to assist in identifying patents from said at least one first database that map to said product, wherein said identified patents each includes at least one claim that potentially reads on said product.

59. The computer program product of claim 53, wherein said processing means comprises:

a patent mapping module to enable the computer to automatically perform a patent mapping function with respect to patents in one of said one or more groups, said patent mapping module comprising:

means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify patents that are within said identified group;

-278-

means for enabling the computer to retrieve information pertaining to said identified patents from at least one of (I) said at least one first database and (II) said at least one second database; and

5 means for enabling the computer to display said retrieved information.

60. The computer program product of claim 59, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group:

10 61. The computer program product of claim 53, wherein said processing means comprises:

a patent citation module to enable the computer to automatically perform a backwards patent citation function with respect to patents in one of said one or more groups, said patent citation module comprising:

15 means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify first patents that are within said identified group;

means for enabling the computer to identify second patents that are cited in said first patents;

20 means for enabling the computer to retrieve information pertaining to said second patents from at least one of (I) said at least one first database and (II) said at least one second database; and

means for enabling the computer to display said retrieved information.

-279-

62. The computer program product of claim 61, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

5 63. The computer program product of claim 53, wherein said processing means comprises:

a patent citation module to enable the computer to automatically perform a forwards patent citation function with respect to patents in one of said one or more groups, said patent citation module comprising:

10 means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify first patents that are within said identified group;

means for enabling the computer to identify second patents that cite said first patents;

15 means for enabling the computer to retrieve information pertaining to said second patents from at least one of (I) said at least one first database and (II) said at least one second database; and

means for enabling the computer to display said retrieved information.

20 64. The computer program product of claim 63, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

65. The computer-program-product of claim 53, wherein said processing means comprises:

-280-

a patent aging module to enable the computer to automatically perform a patent aging function with respect to patents in one of said one or more groups, said patent aging module comprising:

5 means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify patents that are within said identified group;

10 means for enabling the computer to retrieve patent term expiration related information for said identified patents from said at least one first database; and

means for enabling the computer to display patent aging information for said identified patents based on said retrieved patent term expiration related information.

15 66. The computer program product of claim 65, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

67. The computer program product of claim 53, wherein said processing means comprises:

20 a patent bracketing/clustering module to enable the computer to automatically perform a backwards patent bracketing/clustering function with respect to patents in one of said one or more groups, said patent bracketing/clustering module comprising:

means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

25 means for enabling the computer to identify first patents that are within said identified group;

-281-

means for enabling the computer to identify second patents that are cited in said first patents;

means for enabling the computer to retrieve ownership information pertaining to said first and second patents from said at least one second database; and

means for enabling the computer to display a connected graph of said first patents and said second patents, wherein said connected graph depicts said retrieved ownership information.

68. The computer program product of claim 67, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

69. The computer program product of claim 53, wherein said processing means comprises:

a patent bracketing/clustering module to enable the computer to automatically perform a forwards patent bracketing/clustering function with respect to patents in one of said one or more groups, said patent bracketing/clustering module comprising:

means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify first patents that are within said identified group;

means for enabling the computer to identify second patents that cite said first patents;

means for enabling the computer to retrieve ownership information pertaining to said first and second patents from said at least one second database; and

-282-

means for enabling the computer to display a connected graph of said first patents and said second patents, wherein said connected graph depicts said retrieved ownership information.

70. The computer program product of claim 69, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

71. The computer program product of claim 53, wherein said processing means comprises:

an inventor patent count module to enable the computer to automatically perform an inventor patent count function with respect to said patents in one of said one or more groups, said inventor patent count module comprising:

means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify patents that are within said identified group;

means for enabling the computer to retrieve inventorship information pertaining to said identified patents from said at least one second database;

means for enabling the computer to determine top inventors of said identified patents based on said retrieved inventorship information; and

means for enabling the computer to display information pertinent to said top inventors.

72. The computer program product of claim 71, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

-283-

73. The computer program product of claim 53, wherein said processing means comprises:

an inventor employment information module to enable the computer to automatically perform an inventor employment information function with respect to patents in one of said one or more groups, said inventor employment information module comprising:

means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify patents that are within said identified group;

means for enabling the computer to identify inventors of said identified patents by retrieving information pertaining to said identified patents from said at least one second database;

means for enabling the computer to retrieve inventorship information pertaining to said identified inventors from said at least one second database; and

means for enabling the computer to display said retrieved inventorship information.

74. The computer program product of claim 73, wherein said identified group is one of a BOM group, a corporate entity group, an inventor group, and a user-defined group.

75. The computer program product of claim 53, wherein said processing means comprises:

a financial module to enable the computer to automatically perform a financial related function with respect to patents in one of said one or more groups, said financial module comprising:

-284-

means for enabling the computer to receive from an operator information that identifies one of said one or more groups;

means for enabling the computer to identify patents that are within said identified group;

5 means for enabling the computer to retrieve financial information pertaining to said identified patents from said at least one second database; and

means for enabling the computer to display said retrieved financial information.

10 76. The computer program product of claim 53, wherein one of said one or more groups is associated with a person, said computer program logic further comprising:

means for enabling the computer to assist in identifying patents from said at least one first database that name as an inventor said person; and

15 means for enabling the computer to add said identified patents to said one of said one or more groups.

77. The computer program product of claim 53, wherein one of said one or more groups is associated with a second corporate entity, said computer program logic further comprising:

20 means for enabling the computer to assist in identifying patents from said at least one first database that are associated with said second corporate entity; and

means for enabling the computer to add said identified patents to said one of said one or more groups.

25 78. The computer program product of claim 77, wherein said identified patents are owned or licensed by said second corporate entity.

-285-

79. A method of processing and presenting data, comprising the steps of:

(1) maintaining at least one first database of patents;
(2) maintaining at least one second database of non-patent information of interest to a corporate entity;

5 (3) maintaining one or more groups, each of said one or more groups comprising any number of said patents from said at least one first database;

(4) automatically processing said patents in one of said one or more groups in conjunction with non-patent information from said at least one second database; and

10 (5) presenting data generated by said processing in a hyperbolic tree.

80. The method of claim 79, wherein step (4) comprises the step of performing a patent citation tree function or a patent claim tree function.

81. A method of processing and presenting data, comprising the steps of:

15 (1) retrieving patent citation information pertaining to a user-selected patent;

(2) constructing a patent citation hyperbolic tree using said retrieved patent citation information; and

(3) displaying said patent citation hyperbolic tree.

82. A method of processing and presenting data, comprising the steps of:

20 (1) identifying claim dependencies of claims in a user-selected patent;

(2) constructing a patent claims hyperbolic tree for said user-selected patent using said identified claim dependencies; and

(3) displaying said patent claims hyperbolic tree.




FIG.1
PRIOR ART

2/180

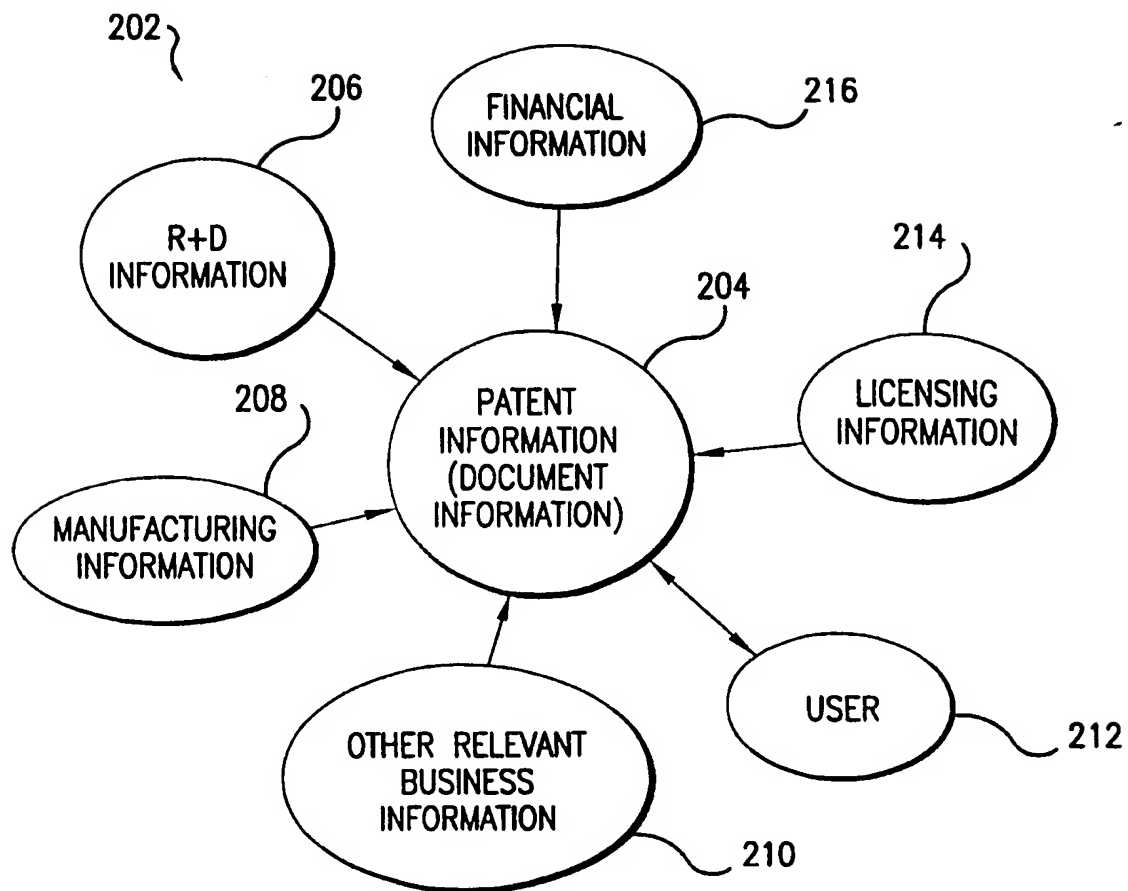


FIG.2





FIG.3



5/180





FIG.5






FIG.7

8/180




FIG.8




FIG.9

10/180

ANALYSIS MODULES 416

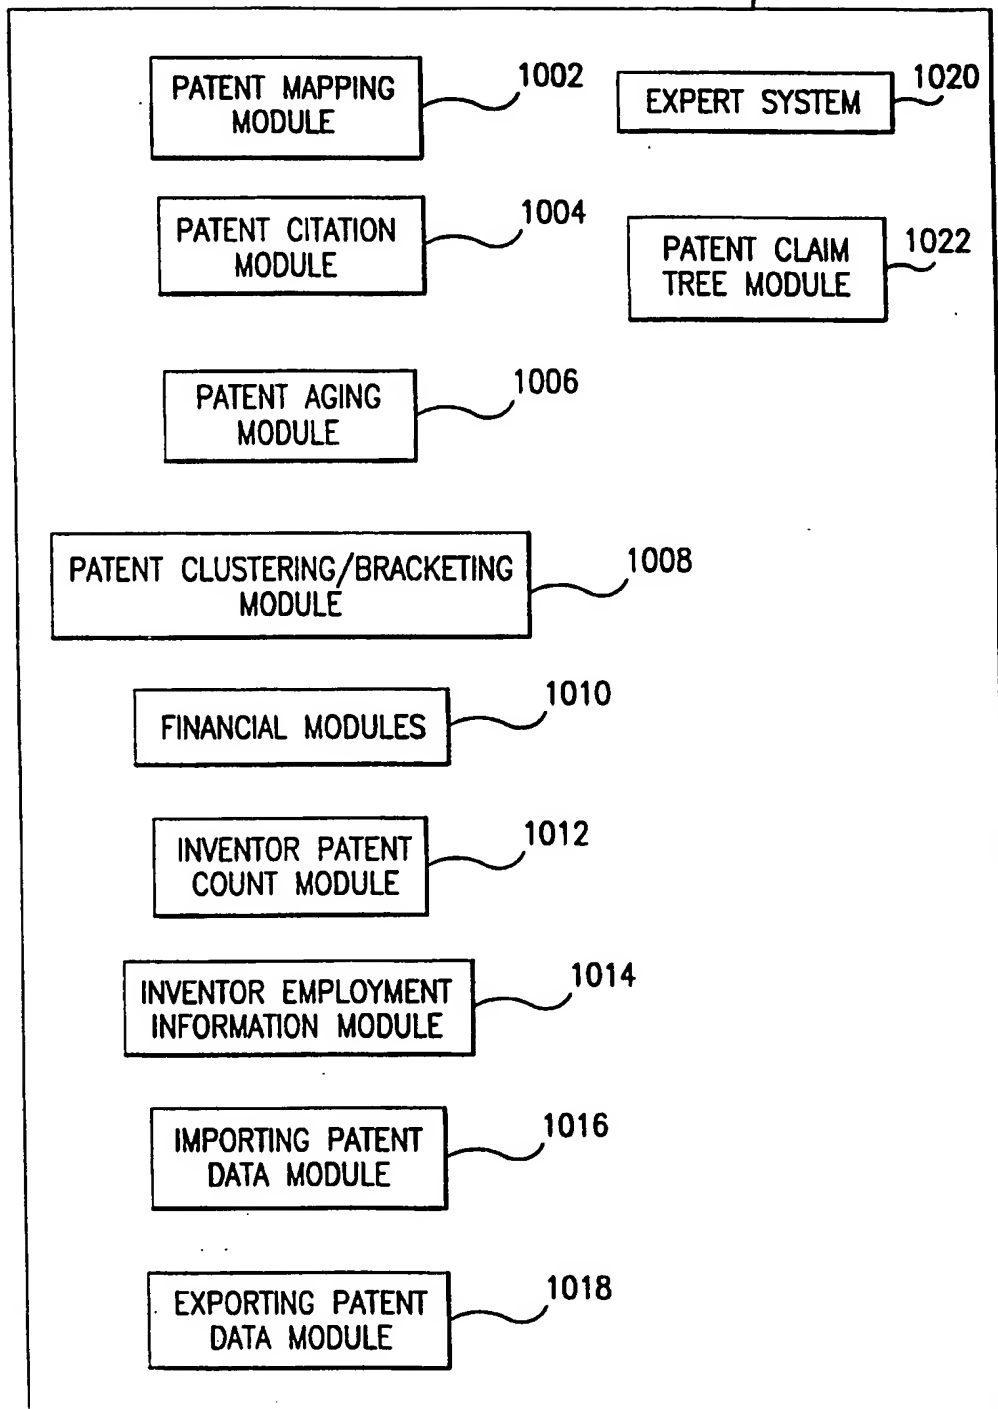


FIG.10




FIG. 11




FIG.12A





FIG. 12B



15/180




FIG.12D

16/180




FIG. 12E

17/180

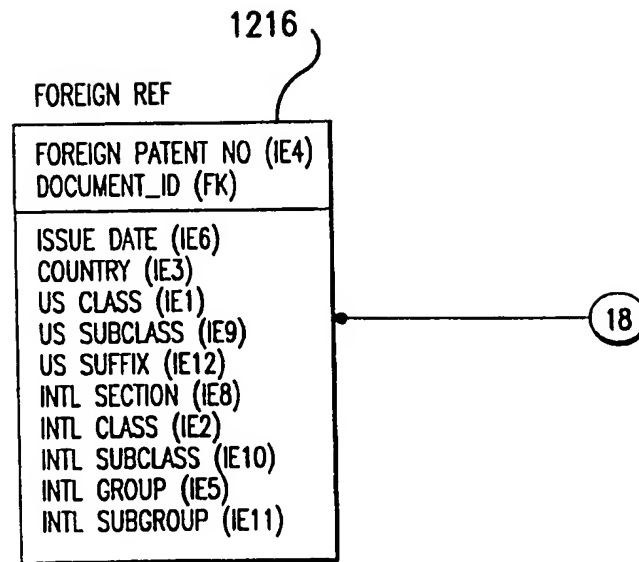


FIG.12F




FIG.12G

19/180




FIG. 12H

20/180

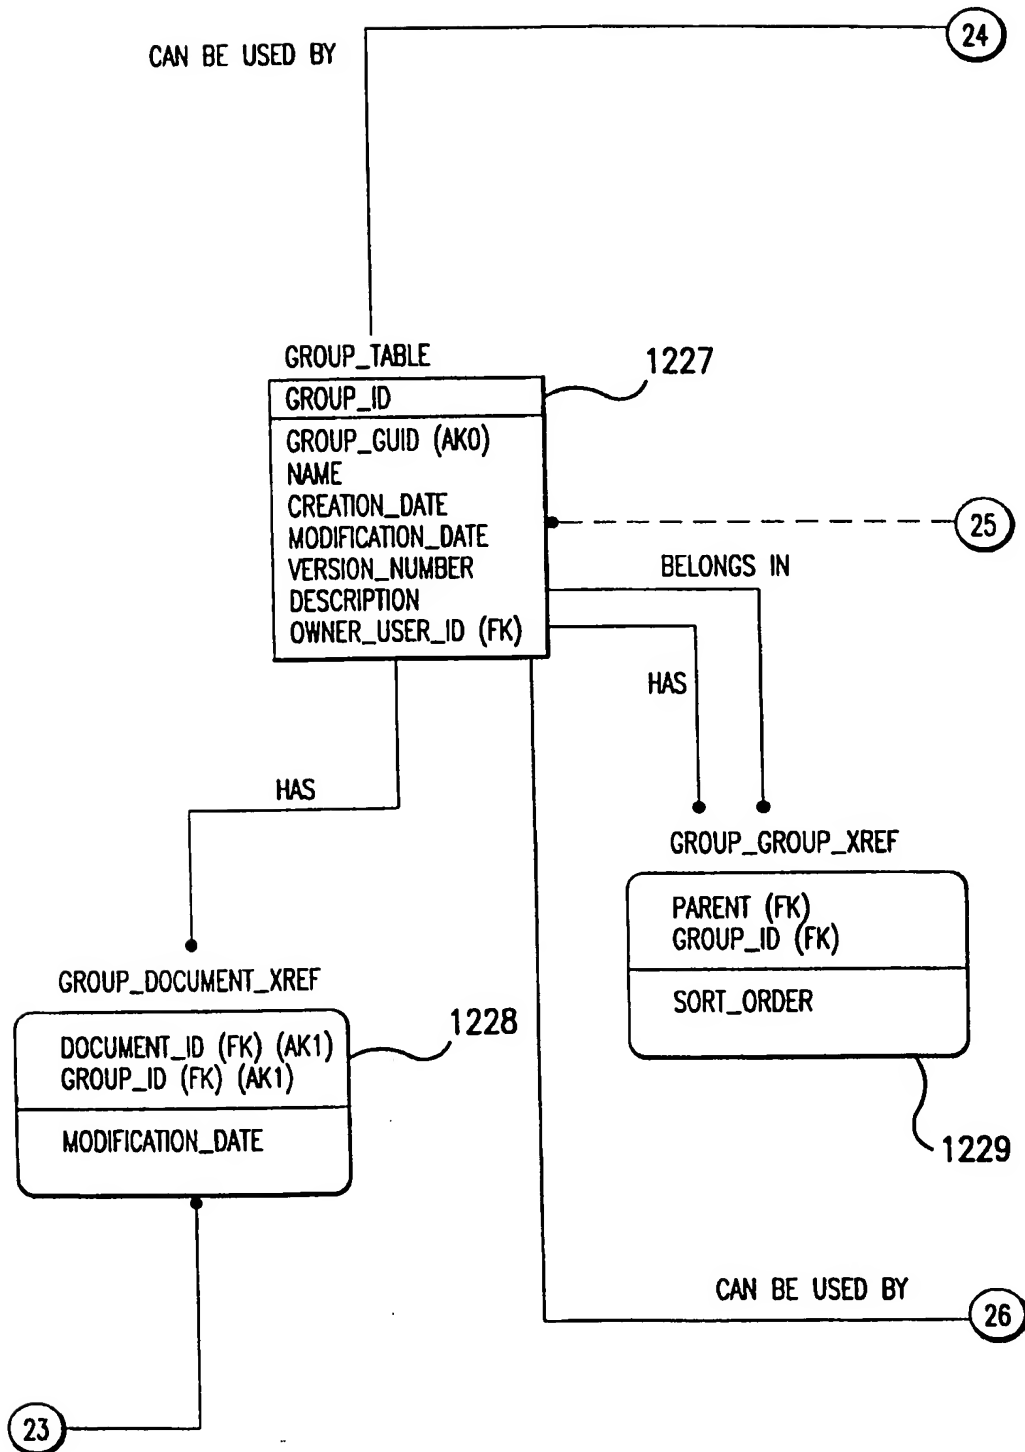


FIG. 12I




FIG.12J




FIG.12K

23/180

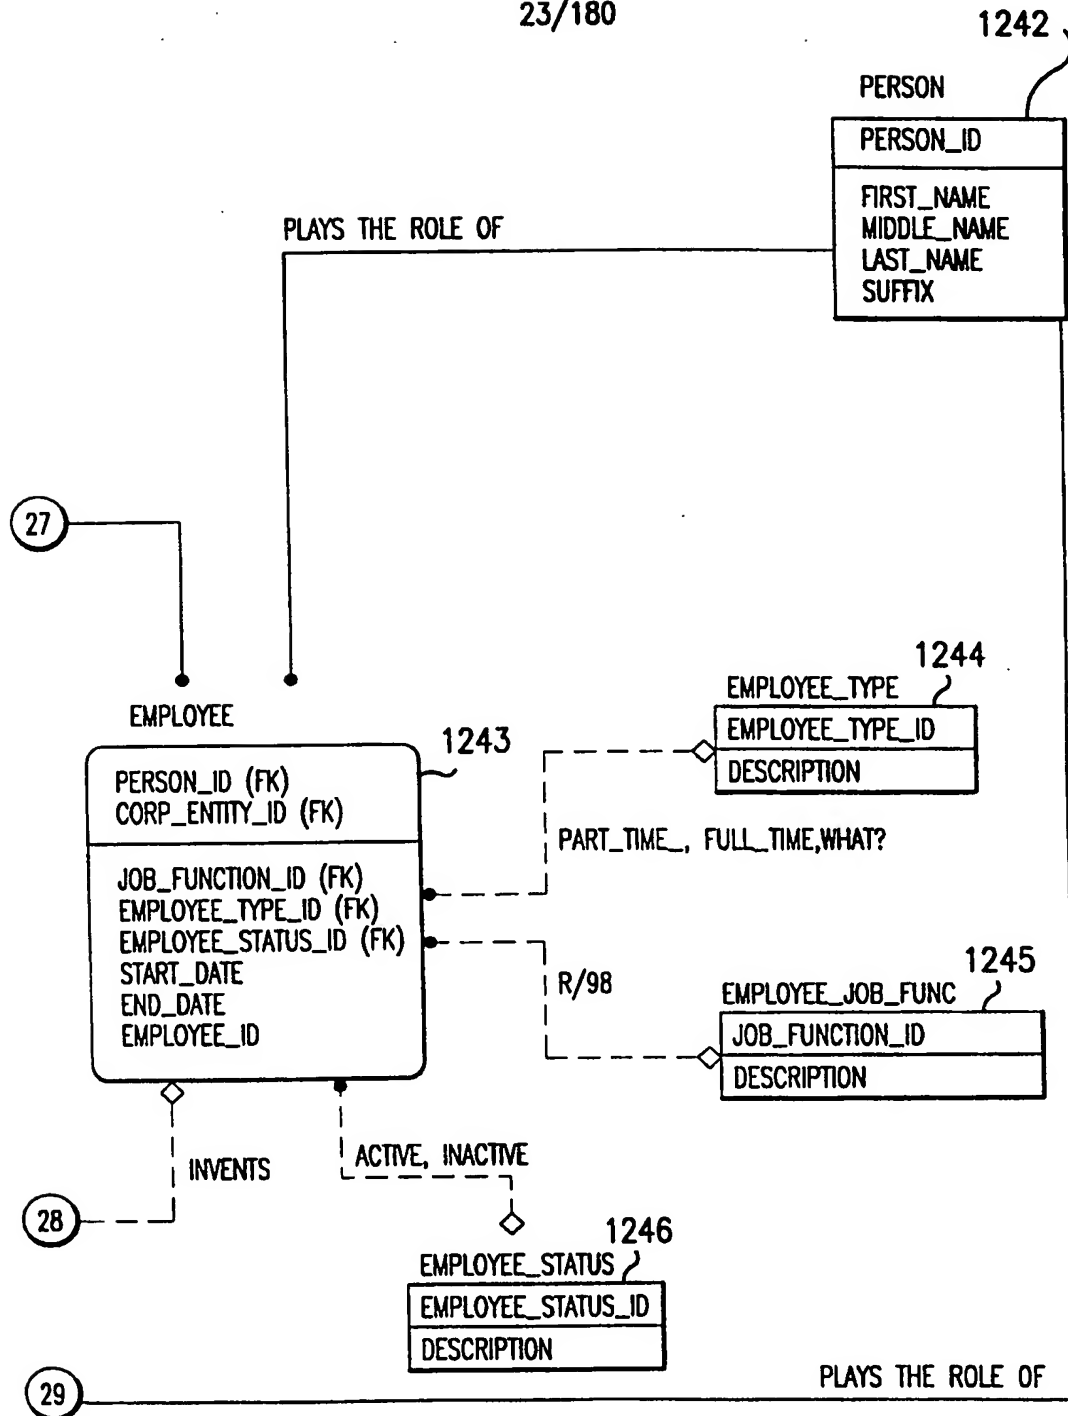



FIG.12L

24/180



PATENT
1222X }

DOCUMENT_ID	DOCUMENT_NUMBER
D1	USP1
D2	USP2
D3	USP3
D4	USP4
D5	USP5
D6	USP6
D7	USP7
D8	USP8
D9	USP9
D10	USP10
D11	USP11
D12	USP12
D13	USP13
D14	USP14

FIG.13

26/180

DEPOSITIONS
1225X }

DOCUMENT_ID	DEPOSITION_ID
D15 D16 D17	DEP1 DEP2 DEP3

FIG.14

PLEADINGS
1224X }

DOCUMENT_ID	PLEADING_ID
D18 D19 D20 D21	PL1 PL2 PL3 PL4

FIG.15

PRIOR_ART
1226X }

DOCUMENT_ID	PRIORART_ID
D22 D23	PA1 PA2

FIG.16

DOCUMENT 27/180
1223K }

DOCUMENT_ID	DOCUMENT_TYPE
D1	PATENT
D2	PATENT
D3	PATENT
D4	PATENT
D5	PATENT
D6	PATENT
D7	PATENT
D8	PATENT
D9	PATENT
D10	PATENT
D11	PATENT
D12	PATENT
D13	PATENT
D14	PATENT
D15	DEPOSITION
D16	DEPOSITION
D17	DEPOSITION
D18	PLEADING
D19	PLEADING
D20	PLEADING
D21	PLEADING
D22	PRIOR ART
D23	PRIOR ART

FIG.17

28/180




FIG.18

29/180

GROUP_TABLE
1227X

GROUP_ID	NAME	OWNER_USER_ID
GRP1	COMPANIES TO ACQUIRE	USER1
GRP2	ABC CORP	USER1
GRP3	XYZ CORP	USER1
GRP4	PQR CORP	USER1
GRP5	PATENTS ABOUT BIKES	USER2
GRP6	LITIGATION	USER3

FIG.19

GROUP_GROUP_XREF
1229X

PARENT	GROUP_ID
GRP1	GRP2
GRP1	GRP3
GRP1	GRP4

2002
2004
2006

FIG.20

30/180

GROUP_DOCUMENT_XREF
1228X

GROUP_ID	DOCUMENT_ID
GRP2	D1
GRP2	D2
GRP2	D3
GRP3	D7
GRP3	D8
GRP4	D11
GRP4	D12
GRP5	D1
GRP5	D2
GRP5	D3
GRP5	D9
GRP6	D22
GRP6	D23
GRP6	D18
GRP6	D21
GRP6	D16
GRP6	D1

FIG.21

2202

31/180

ASSEMBLY

SUBASSEMBLY

PART

BICYCLE

FRAME

SCREW

WHEEL

RIM

TIRE

WHEEL

RIM

TIRE

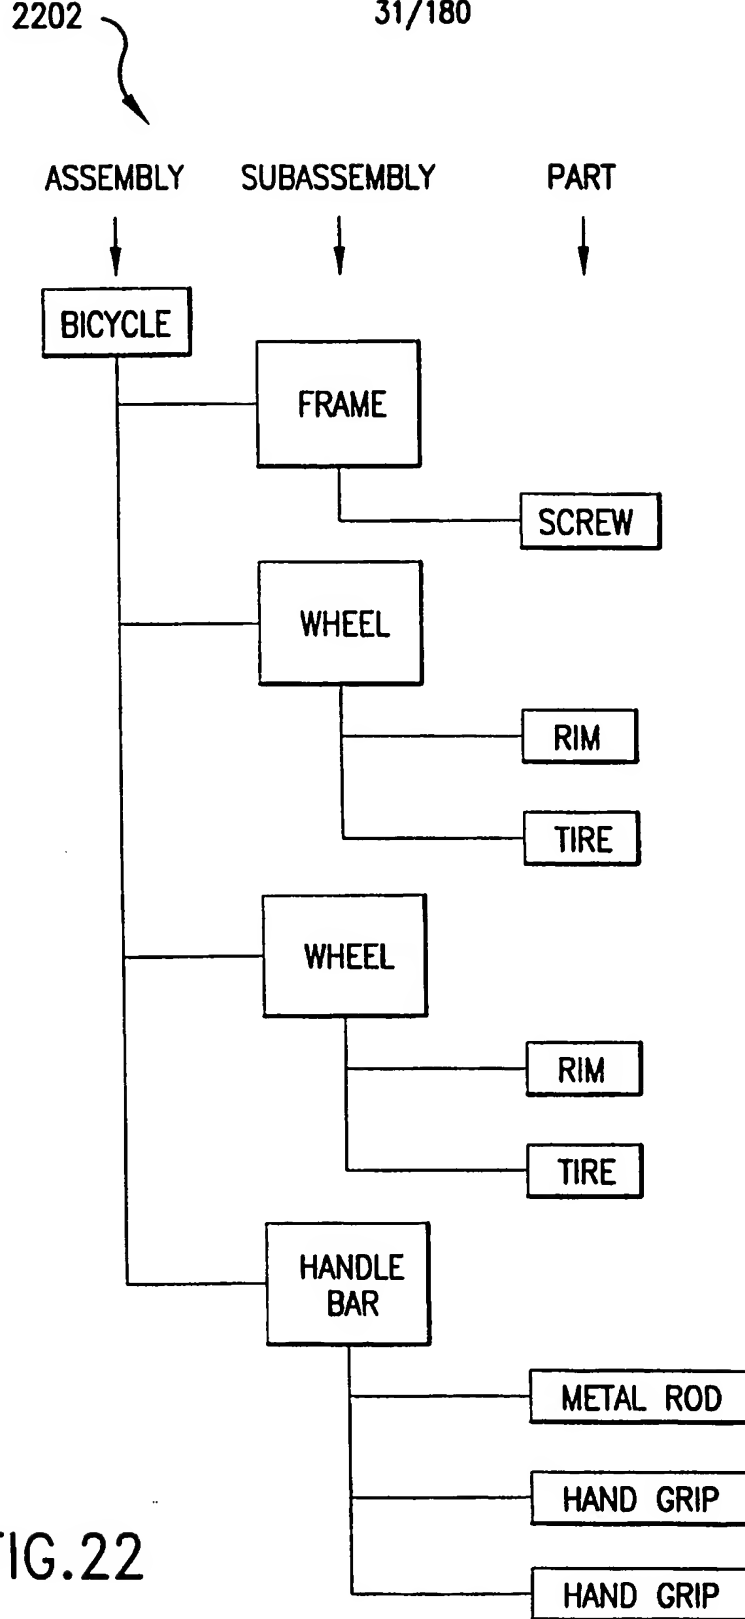
HANDLE
BAR

METAL ROD

HAND GRIP

HAND GRIP

FIG.22



32/180

2302





FIG.23A

33/180




FIG.23B

34/180

BOM
1217X

BOM_ID	NAME
1	BICYCLE
2	BICYCLE FRAME
3	SCREW
4	WHEEL
5	RIM
6	TIRE
7	BICYCLE HANDLE BAR
8	METAL ROD
9	HAND GRIP
10	LAWN MOWER
11	LAWN MOWER FRAME
12	ENGINE
13	LAWN MOWER HANDLE BAR

FIG.24


35/180

BOM_BOM_XREF
1219X

PATENT_BOM_ID	BOM_ID	
1	2	2502
2	3	2504
1	4	2506
4	5	2508
4	6	2510
1	7	2512
7	8	2514
7	9	2516
10	11	2518
11	3	2520
10	12	2522
10	13	2524
13	9	2526


FIG.25

36/180

BOM_PATENT_XREF
1218X

BOM_ID	DOCUMENT_ID
1	D1
1	D2
2	D1
2	D3
3	D4
4	D5
4	D6
5	D5
5	D7
6	D8
7	D9
9	D10
10	D11
10	D12
11	D11
12	D13
12	D14

FIG.26

USER_TABLE
1238X

USER_ID	USER_NAME
USER1	MIKE
USER2	JEN
USER3	KATHRYN

FIG.27

37/180

GROUP_USER_XREF
1237X

GROUP_ID	USER_ID	PERMISSION
GRP1	USER2	CHANGE
GRP2	USER2	NONE

FIG.28

38/180

USER_GROUP
1241X

USER_GROUP_ID	GROUP_NAME
UGRP1	ENGINEERING
UGRP2	PERSONNEL
UGRP3	LEGAL
UGRP4	DAY CARE

FIG.29

USR_GRP_USR_XREF
1239X

USER_GROUP_ID	USER_ID
UGRP2	USER2
UGRP3	USER1
UGRP4	USER3

FIG.30

GRP_USR_GRP_XREF
1240X

GROUP_ID	USER_GROUP_ID	PERMISSION
GRP5	UGRP2	DELETE
GRP5	UGRP3	NONE
GRP6	UGRP3	DELETE
	WORLD	READ

FIG.31

39/180




FIG.32

CORPORATE_ENTITY
1230X

40/180

CORP_ENTITY_ID	CORP_LEVEL_ID	PARENT_CORP_ENTITY_ID
CORP1	PARENT	
CORP2	SUB	CORP1
CORP3	SUB	CORP1
CORP4	SUB	CORP3
CORP5	PARENT	

FIG.33

RELEVANCE
1235X

RELEVANCE_ID	DESCRIPTION
R1	CORE
R2	NON-CORE

FIG.34

PATENT_OWNERSHIP
1234X

OWNERSHIP_ID	DESCRIPTION
P01	OWN
P02	LICENSE

FIG.35

CORP_PATENT_XREF
1233X

CORP_ENTITY_ID	DOCUMENT_ID	OWNERSHIP_ID	RELEVANCE_ID
CORP1	D1	P01	CORE
CORP1	D2	P01	CORE
CORP1	D3	P01	CORE
CORP3	D6	P01	NON-CORE
CORP2	D8	P02	CORE
CORP4	D10	P02	CORE
CORP4	D11	P02	NON-CORE

FIG.36

41/180

PERSON
1242X

PERSON_ID
PERSON1
PERSON2
PERSON3
PERSON4
PERSON5

FIG.37

VALIDATED_INVENTOR
1236X

PERSON_ID	DOCUMENT_ID
PERSON1	D1
PERSON1	D2
PERSON1	D3
PERSON3	D1
PERSON3	D4
PERSON3	D5
PERSON5	D6

FIG.39

EMPLOYEE
1243X

PERSON_ID	CORP_ENTITY_ID	EMPLOYEE_STATUS_ID
PERSON1	CORP1	A
PERSON2	CORP2	A
PERSON4	CORP1	I
PERSON5	CORP3	I

FIG.38


<p>4002 United States Patent [19] Agraliolis et al.</p>	<p>42/180</p>	<p style="text-align: right;">US005574656A</p> <p>[11] Patent Number: 5,574,656 — 4004 [45] Date of Patent: Nov. 12, 1996 — 4006</p>
<p>OTHER PUBLICATIONS</p>		
<p>[54] SYSTEM AND METHOD OF — 4010 AUTOMATICALLY GENERATING CHEMICAL COMPOUNDS WITH DESIRED PROPERTIES 4012 [75] Inventors: Dimitris K. Agraliolis, Exton, Pa.; Roger F. Bone, Bridgewater, N.J.; Francis R. Salemme, Kennett Square, Pa.; Richard M. Soll, Lawrenceville, N.J. [73] Assignee: 3-Dimensional Pharmaceuticals, Inc., Philadelphia, Pa. [21] Appl. No.: 535,822 — 4014 [22] Filed: Sep. 28, 1995 — 4016 4018 Related U.S. Application Data [63] Continuation of Ser. No. 306,915 Sep. 16, 1994, Pat. No. 5,463,564. 4024 [51] Int. Cl. 6 4020 4022 ..G06F 17/50 [52] U.S. Cl.364/500; 364/496; 463/43 [58] Field of Search364/496, 497, 364/500, 4026 499; 436/43, 50, 55; 423/659; 424/2; 935/85-88 [56] References Cited 4028 U.S. PATENT DOCUMENTS 4,939,666 7/1990 Hardman.....436/89 5,240,680 8/1993 Zuckerman et al.422/67 5,270,170 12/1993 Scholtz et al.435/7.37 5,288,514 2/1994 Ellman427/2 5,331,573 7/1994 Balaji et al.364/500 FOREIGN PATENT DOCUMENTS 4040 { 0355628 10/1989 European Pat. Off. . 0355266 2/1990 European Pat. Off. . WO91/19735 12/1991 WIPO. 92/00091 1/1992 WIPO. 93/20242 10/1993 WIPO.</p>	<p>Pabo et al., "Computer-Aided Model — 4008 Building Strategies for Protein Design", <i>Biochemistry</i>, vol. 25, No. 20, pp. 5987-5991 5987-5991, 1986. Saudek et al., "Solution Conformation of Endothelin-1 by HNMR, CD, and Molecular Modeling". <i>International Journal of Peptide Protein Res.</i>, No. 37, pp-174-179 1991. "The Use of Synthetic Peptide Combinatorial Libraries for the Identification of Bioactive Peptides". <i>Peptide Research</i>, vol. 5, No. 6, pp. 351-358, 1992. "Strategies for Indirect Computer-Aided Drug Design", <i>Pharmaceutical Research</i>, vol. 10, No. 4, pp. 475-486. 1993. "Screening Chemically Synthesized Peptide Libraries for Biologically -Relevant Molecules", <i>Organic & Medicinal Chemistry Letters</i>, vol. 3, No. 3, pp. 397- 404, 1993. "Combinatorial Approaches Provide Fresh Leads for Medicinal Chemistry", <i>C&EN</i> Feb. 1994. "Current Trends in Synthetic Peptide and Chemical Diversity Library Design", <i>Genetic Engineering News</i>, pp. 31-32, May 1, 1994. (List continued on next page.) 4030 Primary Examiner - James P. Trammell Assistant Examiner - Kyle J. Choi — 4032 Attorney, Agent, or Firm - Sterne, Kessler, Goldstein & Fox, P.L.L.C. 4034 [57] ABSTRACT A computer based, iterative process for generating chemical entities with defined physical, chemical and/or bioactive properties. During each iteration of the process, (1) a directed diversity chemical library is robotically generated in accordance with robotic synthesis instructions; (2) the compounds in the directed diversity chemical library are analyzed to identify compounds with the desired properties; (3) structure-property data are used to select compounds to be synthesized in the next iteration; and (4) new robotic synthesis instructions are automatically generated to control the synthesis of the directed diversity chemical library for the next iteration.</p>	
		<p>4036 — 2 Claims, 12 Drawing Sheets 4038</p>

FIG. 40

4102

43/180

US005371644A

United States Patent [19]
Hoge et al.

[11] Patent Number: 5,371,644

[45] Date of Patent: *Dec. 6, 1994

[54] SELF IDENTIFYING UNIVERSAL DATA
STORAGE ELEMENT WITH HUMAN
INTELLIGIBLE WRITE PROTECT
MECHANISM

[75] Inventors: David T. Hoge, Westminster;
John C. Owens, Arvada,
both of Colo.;
Michael W. Johnson, Cottage
Grove, Minn.

4104

[73] Assignees: Storage Technology Corporation,
Louisville, Colo.; Minnesota
Mining and Manufacturing
Company, St. Paul, Minn.

4106

[*] Notice: The portion of the term of this
patent subsequent to Aug. 24,
2010 has been disclaimed.

[21] Appl. No.: 115,135

[22] Filed: Sep. 1, 1993

Related U.S. Application Data

[63] Continuation-in-part of Ser. No. 870,578
Apr. 17, 1992, Pat. No. 5,239,437,
which is a continuation-in-part of Ser.
No. 744,456 Aug. 12, 1991, abandoned.

[51] Int. Cl.⁵G11B 15/04

[52] U.S. Cl.360/132;360/133

[58] Field of Search360/132,95,134,131,
360/133;D14/121-123

[56] References Cited

U.S. PATENT DOCUMENTS

D.331,053 11/1992 Zucker et al.D14/115

5,210,671 5/1993 Blackston360/133

5,239,437 8/1993 Hoge et al.360/132

Primary Examiner—Paul M. Dzierzynski

Assistant Examiner—Don Wong

Attorney, Agent, or Firm—Sterne, Kessler,
Goldstein & Fox

[57] ABSTRACT

A mechanism for defining a write protect state of a data storage media enclosed in a housing of a data storage element is disclosed. The mechanism includes an aperture formed in an exterior surface of the housing, wherein the aperture has a first region and a second region. A member is movably secured within the aperture and is manually movable between a first position wherein the member is positioned proximate the aperture first region, and a second position wherein the member is positioned proximate the aperture second region. A first human intelligible mark is placed on the housing proximate the aperture first region. The first human intelligible mark depicts a portion of a predetermined symbol corresponding to one of a write enabled state and a write protect state of the data storage media. A second human intelligible mark is placed on the member. The second human intelligible mark depicts a remaining portion of the symbol. The second human intelligible mark is alignable with the first human intelligible mark by moving the member into the first position. When the first and second human intelligible marks are aligned, the first and second human intelligible marks depict an unbroken image of the symbol, thereby indicating that the data storage media is in the one of the write enabled state and the write protect state. When the first and second human intelligible marks are not aligned, the first and second human intelligible marks depict a broken image of the symbol, thereby indicating that the data storage media is in another of the write enabled state and the write protect state of the data storage media.

7 Claims, 4 Drawing Sheets




FIG. 41

44/180

United States Patent [19]
Morel et al.4202US005302779A
[11] Patent Number: 5,302,779
[45] Date of Patent: Apr. 12, 1994[54] PLUG FOR FIXING IN AN IMPERMEABLE
MANNER AN ELECTRIC CABLE TO AN
OPENING AND CABLE PROTECTION SLEEVE
COMPRISING SUCH PLUGS[75] Inventors: Jacques Morel; Didier Morel, both of
Chateauf-neuf-en-Thymerois, France[73] Assignee: Etablissements Morel-Ateliers
Electromechaniques de Favieres,
Chateauf-neuf-en-Thymerois, France

[21] Appl. No.: 938,138

[22] PCT Filed: Feb. 20, 1991

[86] PCT No.: PCT/FR91/00135

4204 § 371 Date: Dec. 10, 1992

4206 § 102(e) Date: Dec. 10, 1992

[87] PCT Pub. No.: WO92/15138

PCT Pub. Date: Sep. 3, 1992

[51] Int. Cl.⁵ H02G 15/08[52] U.S. Cl. 174/92; 174/76
174/77 R; 174/93[58] Field of Search 174/92, 93, 76, 77 R,
174/82, 65 G, 41

[56] References Cited

U.S. PATENT DOCUMENTS

3,769,443 10/1973 Pierzchala et al. ... 174/76 X
 3,935,373 1/1976 Smith et al. 174/77 R
 4,262,167 4/1981 Bossard et al. 174/92
 4,666,240 5/1987 Caron et al. 174/93
 4,704,499 11/1987 Faust 174/92
 4,752,653 6/1988 Bachel et al. 174/93
 4,861,946 8/1989 Pichler et al. 174/92
 4,933,512 6/1990 Nimiya et al. 174/92

FOREIGN PATENT DOCUMENTS

0023099 1/1981 European Pat. Off. .
 0073748 3/1983 European Pat. Off. .
 0098765 1/1984 European Pat. Off. .
 0206854 12/1986 European Pat. Off. .
 0246113 11/1987 European Pat. Off. .
 419820 10/1925 Fed. Rep. of Germany.
 2826584 10/1979 Fed. Rep. of Germany.
 2466887 4/1981 France .
 2596215 9/1987 France .
 2193605 2/1988 United Kingdom .

Primary Examiner— Morris H. Nimmo
 Attorney, Agent, or Firm— Young & Thompson

[57] ABSTRACT

A plug that makes it possible to fix in an impermeable manner an electric or telephone cable (2) to a circular opening (4) for passage of the cable. It comprises a hollow body (5) constituted by two shells intended to be fixed around the cable (2) and comprising at least two jaws (6, 7) which are retained radially in the body (5). Each jaw comprises a central opening (10, 11) for passage of the cable (2) and is supported radially on this cable at the time of fixing by clamping of the two shells to one another, the two jaws defining between them a chamber (12) intended to be filled with a sealing material. The hollow body (5) has structure for fixing it in a removable and impermeable manner to the circular opening (4) for passage of the cable.

24 Claims, 6 Drawing Sheets




FIG. 42

45/180

United States Patent [19]
Grief et al.4302[11] E Patent Number: RE34583E 4304 Re. 34,583
[45] Reissued Date of Patent: Apr. 12, 1994[54] METHOD OF FORMING A
CONFIGURATION OF INTERCONNECTIONS
ON A SEMICONDUCTOR DEVICE HAVING A
HIGH INTEGRATION DENSITY

FOREIGN PATENT DOCUMENTS 4306

0165085 12/1985 European Pat. Off. .

Primary Examiner— William A. Powell
Attorney, Agent, or Firm— Norman N. Spain

[57] ABSTRACT

[75] Inventors: Malcom K. Grief; Trung T. Doan,
both of Boise, Id.; Hendrikus
J.W. Van Houtum; Josephus M.H.L.
Van Laarhoven, both of Eindhoven,
Netherlands[73] Assignee: U.S. Philips Corporation, New York,
N.Y.

[21] Appl. No.: 880,860

[22] Filed: May 11, 1992

Related U.S. Patent Documents

Reissue of:

[64] Patent No.: 4,936,950
4308 Issued: Jun. 26, 1990
Appl. No.: 339,029
Filed: Apr. 14, 1989

4310

[30] Foreign Application Priority Data

Apr. 22, 1988 [FR] France.....88 05391

[51] Int. Cl.⁵.....H01L 21/306; B44C 1/22;
C23F 1/00; C03C 15/00[52] U.S. Cl.156/643; 156/644;
156/646; 156/653; 156/657; 156/659.1;
156/662; 257/750; 437/192; 437/203[58] Field of Search156/643, 644, 646,
653, 156/656, 657, 659.1, 662; 437/203, 246,
194; 357/71

[56] References Cited

U.S. PATENT DOCUMENTS

4,184,909 1/1980 Chang et al.437/192 X
4,361,599 11/1982 Wourms.....437/192 X
4,484,978 11/1984 Keyser.....156/653 X
4,536,249 8/1985 Rhodes.....156/653 X
4,624,864 11/1986 Hartmann.....437/193 X
4,640,738 2/1987 Fredericks et al.156/656
4,981,550 1/1991 Hultemomn et al. ...437/194 X

A method of the kind consisting in that a contact is obtained with an active zone (11) carried by a semiconductor substrate (10) by means of conductive contact studs (18a) located in the contact openings (16c) of an isolating layer (12) and in that then a metallic configuration of interconnections (22) is formed establishing the conductive connection with the conductive contact studs (18a). A separation layer (13) is provided between the isolating layer (12) and the conductive layer (18), which can be eliminated selectively with respect to the isolating layer (12). Thus, the isolating layer (12) retains its original flatness and the conductive contact studs (18a) have an upper level (20) exceeding slightly the level (21) of the isolating layer (12), thus favoring the contact between these contact studs (18a) and the metallic configuration of interconnections (22). Application in microcircuits having a high integration density.





FIG.43

United States Patent [19] Rivette et al.	4402 46/180	US005623681A [11] Patent Number: 5,623,681 [45] Date of Patent: Apr. 22, 1997																																							
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>[54] METHOD AND APPARATUS FOR SYNCHRONIZING DISPLAYING AND MANIPULATING TEXT AND IMAGE DOCUMENTS</p> <p>[75] Inventors: Kevin G. Rivette, Palo Alto; Michael P. Florio, Atherton; Adam Jackson, Sunnyvale; Don Ahn, Daly City; Irving S. Rappaport, Palo Alto; Deborah Kurata, Pleasanton, all of Calif.</p> <p>[73] Assignee: Waverly Holdings, Inc., Palo Alto, Calif.</p> <p>[21] Appl. No.: 155,752</p> <p>[22] Filed: Nov. 19, 1993</p> <p>[51] Int. Cl.⁶ G06F 17/00</p> <p>[52] U.S. Cl. 395/788; 395/784; 395/787; 395/341</p> <p>[58] Field of Search 395/144-149, 395/155, 157, 158, 160, 161, 788, 784, 783, 787, 341; 364/401-408, 419; 382/135, 137, 140</p> <p>[56] References Cited U.S. PATENT DOCUMENTS</p> <table style="width: 100%; border-collapse: collapse;"> <tr><td style="width: 30%;">Re. 32,632</td><td style="width: 40%;">3/1988 Atkinson.....</td><td style="width: 30%;">345/165</td></tr> <tr><td>4,205,780</td><td>6/1980 Burns et al.</td><td>235/54</td></tr> <tr><td>4,270,182</td><td>5/1981 Asija</td><td>382/192</td></tr> <tr><td>4,486,857</td><td>12/1984 Heckel</td><td>395/100</td></tr> <tr><td>4,533,910</td><td>8/1985 Sukonick et al.</td><td>345/118</td></tr> <tr><td>4,555,775</td><td>11/1985 Pike.....</td><td>345/158</td></tr> <tr><td>4,622,545</td><td>11/1986 Atkinson.....</td><td>345/191</td></tr> <tr><td>4,736,308</td><td>4/1988 Heckel</td><td>395/144</td></tr> <tr><td>4,748,618</td><td>5/1988 Brown et al.</td><td>370/191</td></tr> <tr><td>4,772,882</td><td>9/1988 Mical</td><td>345/146</td></tr> <tr><td>4,785,408</td><td>11/1988 Britton et al.</td><td>395/2.79</td></tr> <tr><td>4,788,538</td><td>11/1988 Klien et al.</td><td>395/118</td></tr> <tr><td>4,812,834</td><td>3/1989 Wells</td><td>364/500</td></tr> </table> <p style="text-align: center;">(List continued on next page.)</p> <p style="text-align: center;">FOREIGN PATENT DOCUMENTS</p> <p>0239884 10/1987 European Pat. Off. .</p> </div> <div style="width: 50%;"> <p style="text-align: center;">OTHER PUBLICATIONS</p> <p>Alexander, Visualizing cleared-off desktops, <i>Computerworld</i> May 6, 1991, p.20 Hiroshi Ishii, Kouho Arito, <i>Clearface: Translucent Multiuser Interface for Team Workstation</i>, <i>ECSCW</i>, Sep. 1991, pp. 6-10.</p> <p style="text-align: center;">(List continued on next page)</p> <p>Primary Examiner - Almis R. Jankus Assistant Examiner - Joseph R. Burwell Attorney, Agent or Firm - Sterne, Kessler, Goldstein & Fox, P.L.L.C.</p> <p>[57] ABSTRACT</p> <p>The present invention provides a method and apparatus for extracting, synchronizing, displaying, and manipulating text and image documents in electronic form for display. In the Preferred embodiment of the present invention, text and image files for documents, such as for example patent documents, are initially stored on separate magnetic tape media. These data files are extracted from the respective tapes and placed onto a faster medium, such as a hard disk drive. The text and image files are synchronized to produce Equivalent Files using heuristic algorithms to create an approximate equivalence relationship between the text and the image files. The Equivalent Files and image files residing on for example, a hard disk drive or compact disk (CD), are coupled as a resource to a computer display system. The graphic user interface of the present invention permits the user to display, manipulate, and edit the Equivalent File created using the process of the present invention, and to simultaneously view the image file on the display. Using the graphic user interface of the present invention, a user may create libraries of patent Equivalent Files and image files or documents of another type, as well as open cases to include a plurality of different patents or other documents. The Equivalent File may be selectively viewed on the display in an equivalent window. Simultaneous with the viewing of the Equivalent File within the equivalent window, the user may view any portion of an image file within one or more image windows on the display. A variety of other features and functions are provided by the present invention for the manipulation and display of patent documents and other documents on the user interface.</p> </div> </div>			Re. 32,632	3/1988 Atkinson.....	345/165	4,205,780	6/1980 Burns et al.	235/54	4,270,182	5/1981 Asija	382/192	4,486,857	12/1984 Heckel	395/100	4,533,910	8/1985 Sukonick et al.	345/118	4,555,775	11/1985 Pike.....	345/158	4,622,545	11/1986 Atkinson.....	345/191	4,736,308	4/1988 Heckel	395/144	4,748,618	5/1988 Brown et al.	370/191	4,772,882	9/1988 Mical	345/146	4,785,408	11/1988 Britton et al.	395/2.79	4,788,538	11/1988 Klien et al.	395/118	4,812,834	3/1989 Wells	364/500
Re. 32,632	3/1988 Atkinson.....	345/165																																							
4,205,780	6/1980 Burns et al.	235/54																																							
4,270,182	5/1981 Asija	382/192																																							
4,486,857	12/1984 Heckel	395/100																																							
4,533,910	8/1985 Sukonick et al.	345/118																																							
4,555,775	11/1985 Pike.....	345/158																																							
4,622,545	11/1986 Atkinson.....	345/191																																							
4,736,308	4/1988 Heckel	395/144																																							
4,748,618	5/1988 Brown et al.	370/191																																							
4,772,882	9/1988 Mical	345/146																																							
4,785,408	11/1988 Britton et al.	395/2.79																																							
4,788,538	11/1988 Klien et al.	395/118																																							
4,812,834	3/1989 Wells	364/500																																							
																																									
FIG. 44A																																									

51 Claims, 71 Drawing Sheets

47/180

US005623679A

United States Patent [19]
Rivette et al.

[11] Patent Number: 5,623,679
[45] Date of Patent: Apr. 22, 1997

[54] SYSTEM AND METHOD FOR CREATING AND
MANIPULATING NOTES EACH CONTAINING
MULTIPLE SUB-NOTES, AND LINKING THE
SUB-NOTES TO PORTIONS OF DATA OBJECTS

[75] Inventors: Kevin G. Rivette, Palo Alto; Michael
P. Florio, Atherton; Adam Jackson,
Belmont; Don Ahn, Daly City; Irving
S. Rappaport, Palo Alto; Deborah
Kurata, Pleasanton, all of Calif.

4410

[73] Assignee: Waverly Holdings, Inc., Palo Alto, Calif.

[21] Appl. No.: 423,676

[22] Filed: Apr. 18, 1995

4490 Related U.S. Application Data

[63] Continuation-in-part of Ser. No. 341,129,
Nov. 18, 1994, which is a continuation-in-
part of Ser. No. 155,752, Nov. 19, 1993.

[51] Int. Cl. ⁶ G06F 17/00

[52] U.S. Cl. 395/773; 395/784; 395/788;

[58] Field of Search 395/144-149,
395/155, 157, 160, 161, 773, 784, 788,
364/419.19

[56] References Cited
U.S. PATENT DOCUMENTS

Re. 32,632 3/1988 Atkinson 345/165
4,205,780 6/1980 Burns et al. 235/454
4,270,182 5/1981 Asija 382/192
4,486,857 12/1984 Heckel 395/100
4,533,910 8/1985 Sukonick et al. 345/118
4,555,775 11/1985 Pike 395/158
4,622,545 11/1986 Atkinson 345/191
4,736,308 4/1988 Heckel 395/144

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

0239884 10/1987 European Pat. Off. .

OTHER PUBLICATIONS

Collin et al., "InterNote: Extending a
Hypermedia Framework to support
Annotative Collaboration", Hypertext '89
Proceedings, pp. 365-378. Nov. 1989.
Fish et al., "Quilt: a collaborative tool for
Cooperative writing", Conf. on Information
Systems '88, pp. 30-37. Jan. 1988.
Thompson et al., FullWrite Professional-A
User's Guide, pp. 99-122. Jan. 1988

(List continued on next page)

[57] ABSTRACT

The present invention provides a method and
apparatus for extracting, synchronizing, displaying,
and manipulating text and image documents in
electronic form for display. In the Preferred
embodiment of the present invention, text and
image files for documents, such as for example
patent documents, are initially stored on separate
magnetic tape media. These data files are extracted
from the respective tapes and placed onto a faster
medium, such as a hard disk drive. The text and
image files are synchronized to produce Equivalent
Files using heuristic algorithms to create an
approximate equivalence relationship between the
text and the image files. The Equivalent Files and
image files residing on for example, a hard disk
drive or compact disk (CD), are coupled as a
resource to a computer display system. The graphic
user interface of the present invention permits
the user to display, manipulate, and edit the
Equivalent File created using the process of the
present invention, and to simultaneously view the
image file on the display. Using the graphic user
interface of the present invention, a user may
create libraries of patent Equivalent Files and
image files or documents of another type, as
well as open cases to include a plurality of different
patents or other documents. The Equivalent File
may be selectively viewed on the display in an
in an equivalent window. Simultaneous with the
viewing of the Equivalent File within the equivalent
window, the user may view any portion of an image
file within one or more image windows
on the display. A variety of other features and
functions are provided by the present invention
for the manipulation and display of patent
documents and other documents on the
user interface.

10 Claims, 84 Drawing Sheets




FIG. 44B




FIG.45




FIG. 46

50/180




FIG.47




FIG. 48

52/180




FIG. 49




FIG.50




FIG. 51





FIG.52

56/180

5302


SMARTPATENTS

PATENT SEARCH

5304 SCOPE OF SEARCH: ALL U.S. PATENTS ↓ 5306

PATENT NUMBER:

TITLE: 5306 5308

INVENTOR: 5310

ASSIGNEE: 5312

CLASS: 5316 5314

DATE OF ISSUE: ☒ BEFORE ↓ 5318 (MM/dd/yyyy)

5315 5320 ☐ BETWEEN AND (MM/dd/yyyy)

ABSTRACT: 5322 5324 5326

FULL PATENT TEXT: 5328

FIELDS TO INCLUDE IN LIST OF SEARCH RESULTS (PATENT NUMBER IS ALWAYS INCLUDED)

<input checked="" type="checkbox"/> TITLE	<input type="checkbox"/> DATE OF ISSUE 5330
<input type="checkbox"/> INVENTOR	<input type="checkbox"/> ASSIGNEE

ORDER RESULTS BY SCORE ↓ 5332

5334

SEARCH

5336

HELP WITH THIS SCREEN

FIG.53

57/180

5402

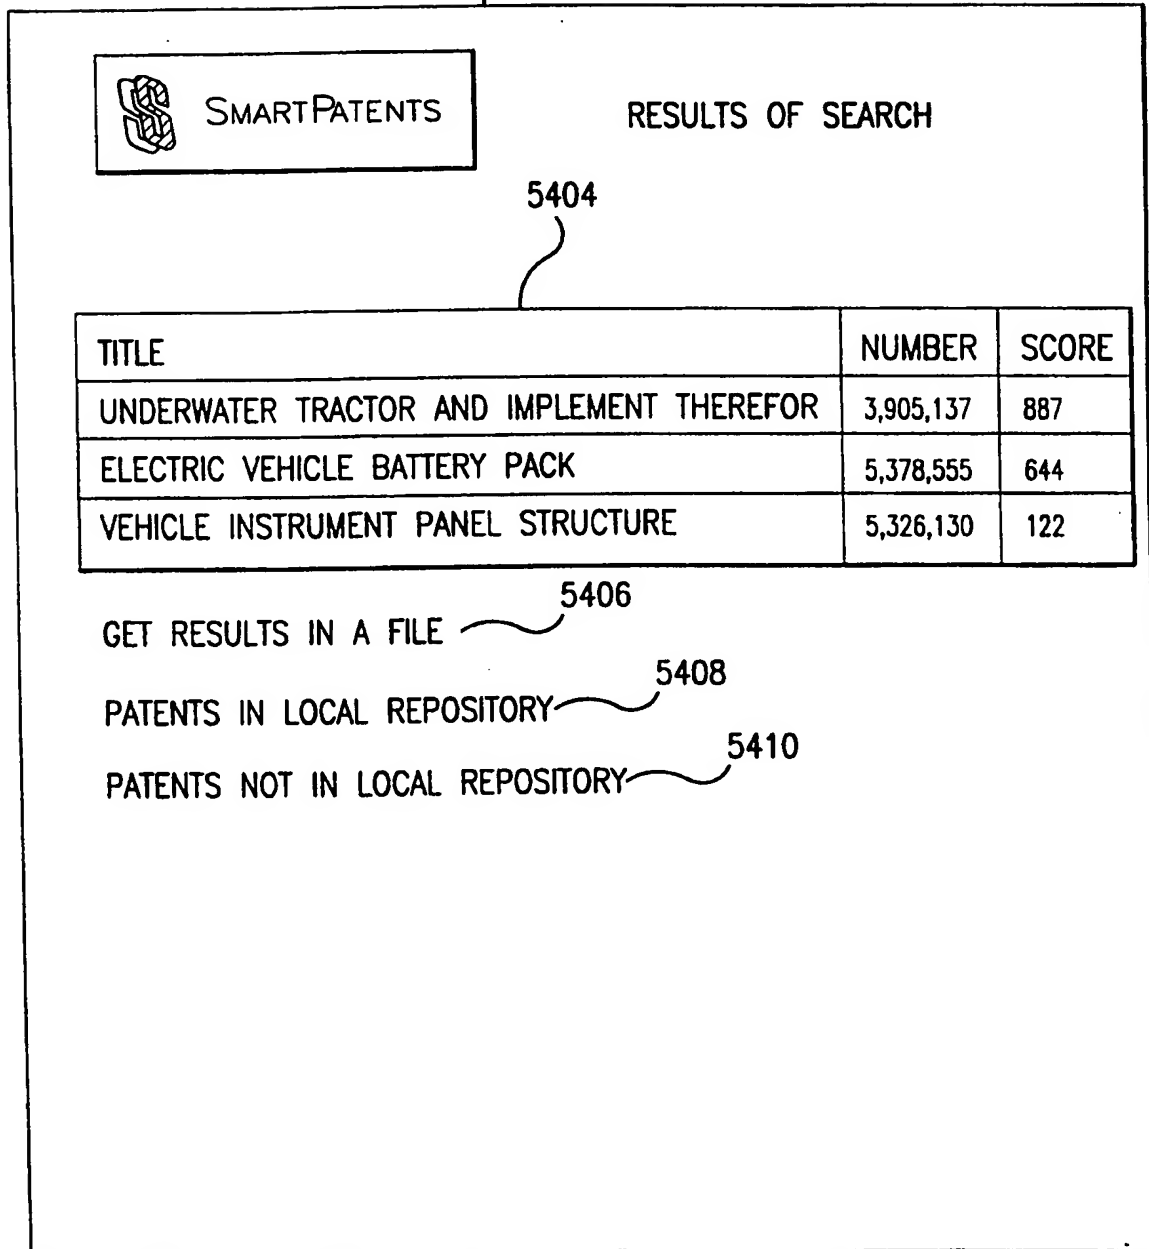


FIG.54

58/180

SMART PATENTS		
LOOKUP	TEXT	IMAGE
5504	5506	5508
BIBLIOGRAPHY ~ 5510		
REFERENCES ~ 5512		
CLAIMS ~ 5514		
UNITED STATES PATENT [19]		
GEE		
[11]	PATENT NUMBER:	3,905,137
[45]	DATE OF PATENT:	SEP. 16, 1975

[54]	UNDERWATER TRACTOR AND IMPLEMENT THEREFOR	
[75]	INVENTOR: JAMES E. GEE, WASHINGTON, IL	
[73]	ASSIGNEE: CATERPILLAR TRACTOR COMPANY, PEORIA, IL	
[21]	APPL. NO.: 444,317	
[22]	FILED: FEB. 21, 1974	
RELATED U.S. APPLICATION DATA		
[51]	INT. CL. 2E02F 3/88	
[52]	U.S. CL. 37/56; 37/66; 37/71; 61/69; 180/1.H	
[58]	FIELD OF SEARCH 37/56, 71, 66, 54; 180/1 H; 61/69 R, 69 A; 115/1 R	
[56]	REFERENCES CITED	
	U.S. PATENT DOCUMENTS	
	658,922 10/1900 OVERSTROM.....	37/56
	813,935 2/1906 AVERY, JR.	37/56
	3,314,174 4/1967 HAGGARD.....	37/54
	3,680,521 8/1972 MURAKI ET AL.	115/1R
	3,683,521 8/1972 SLOAN ET AL.	37/56
	3,706,142 12/1972 BRUNNER.....	37/56
	3,738,029 6/1973 HARMON.....	37/66
	3,757,438 9/1973 WATASE.....	37/56
	3,774,323 11/1973 VAUGHN.....	37/71 X
	FOREIGN PATENT DOCUMENTS	
	216,537 7/1968 SOVIET UNION.....	37/66
	841,472 5/1970 CANADA.....	37/66
PRIMARY EXAMINER--CLIFFORD D. CROWDER		
ATTORNEY, AGENT, OR FIRM--PHILLIPS, MOORE, WEISSENBARGER LEMPIO & STRABALA		
[57]	ABSTRACT	
A TRACTOR INCLUDES AN INTERNAL COMBUSTION ENGINE AND IS		

FIG.55

59/180

United States Patent [19]
Gee5602[11] 3
[45] Sept[54] UNDERWATER TRACTOR AND
IMPLEMENT THEREOF

841,472 5/1970 Canada

[75] Inventor: James E. Gee, Washington, Ill.

Primary Examiner Clifford D. Crowde
Attorney, Agent or Firm Phillips, Moo
Weissenberger Lempio & Strubolo[73] Assignee: Caterpillar Tractor Company,
Peoria, Ill.

[22] Filed: Feb. 21, 1974

[21] Appl. No.: 444,317

[57] ABSTRACT

[52] U.S. Cl.37/56; 37/66; 37/71;
61/69 R; 180/1 H[51] Int. Cl.².....E02F 3/68[58] Field of Search37/56,71,66,54;
180/1 H; 61/69 R 69 A; 115/1 R[56] References Cited
UNITED STATES PATENTS

658,922	10/1900	Overston.....	37/56
813,935	2/1906	Avery, Jr.	37/56
3,314,174	4/1967	Haggard	37/54
3,680,521	8/1972	Minski et al. ..	15/1 R
3,683,521	8/1972	Sloan et al. ...	27/56
3,706,142	12/1972	Brunner	37/56
3,738,029	6/1973	Harper.....	37/66
3,757,436	9/1973	Watusi.....	37/56 UX
3,724,323	11/1973	Vaughn.....	37/71 X

A tractor includes an internal combustion is adapted for underwater operation , A outlet means extending to above the s water for supplying air to the engine a combustion gas therefrom. A dredging the form of a blade and auger, directs the floor to adjacent one end of a con the vehicle is transported over the floor associated with the conduit means trans lerial therethrough and from the other e duit means, which is positioned above l the water. The dredging implement may take the form of a pair of wing member upwardly and inwardly in the rearward d nect material to adjacent the one end e means . Means remote from the vehicle able on shore, are included for opera merged tractor.

4 Claims, 3 Drawing Figure


FOREIGN PATENTS OR APPLICATIONS
216,537 7/1968 U.S.S.R.37/66

FIG.56

60/180

Search

Search What?

☒ Full Text Index

☐ Repository

☐ Current Group

☐ Current Patent

Bib Fields

Patent It:

Title:

Assignee:

Issue Date: ▾ And

Keywords:

Inventor Name

Original U.S. Class

Class:

Subclass:

Suffix:

First:

Last:


Text

Abstract:

Full Text:

5702

FIG.57



5802

-16.58

63/180

Phone Handset <BOM level>						
Part Descr	Part No	Part Status	Patent No	Ownership	Core Status	Ranking
cord	N 33a	production	5003922	owned	Core	10
top casing	300be #3	production	5103944	licensed	Core	10
bottom casing	300sal	production	4739181	owned	non-core	3
Phone Innards <BOM level>						
Part Descr	Part No	Part Status	Patent No	Ownership	Core Status	Ranking
doohickey1	A3383 F	production	5003922	owned	Core	10
doohick	B 0983893	production	5103944	licensed	Core	10
dongle	C 383a	production	4739181	owned	non-core	3

FIG.60

64/180

6118

6102

Patent Citation Report					
Digital Video Recording System 6104					
Src Pat	Src Patent Title	Citing Pat	Assignee	Title	Date
5327235	Video conversions of vide	5345264	Sanyo Electric Co.	Video signal process	19940906
		5363264	The United States	Versatile digital recon	19941108
		5371602	Hitachi, Ltd.	Picture data recording	19941206
		5374958	Sonora	Image compression b	19941220
		5377014	AT&T Corp.	Apparatus and metho	19941227
		5381274	Sonora	Apparatus and metho	19950110
		5438459	Hitachi, Ltd.	Method of processing	19950801
6106 { 6108 6110 6112 6114 6116 }					
Src Pat	Src Patent Title	Citing Pat	Assignee	Title	Date
5359428	Digital video recorder ope	5379072	Sonora	Digital video signal re	19950103
		5381176	Sonora	Miniaturized video com	19950110
		5386231	Sonora	Video camera	19950131
		5387932	Sonora	Video camera capabl	19950207
		5390158	Sonora	Method for recording	19950214
		5521637	Sonora	Solid stste image picl	19960528

6118

FIG. 61

65/180

Source Pat No		Source Pat Title			
		Citing Pat	Assignee	Title	IssueDate
		Patent No.	Assignee	Title	Date
		Patent No.	Assignee	Title	Date
					Ownership
					Ownership
					Ownership
Source Pat No		Source Pat Title			
		Citing Pat	Assignee	Title	IssueDate
		Patent No.	Assignee	Title	Date
		Patent No.	Assignee	Title	Date
					Ownership
					Ownership
					Ownership

FIG.62

<BOM-level>		Handset SubAssembly				
Source Patent	Src Pat Title	Citing Pat	Assignee	Title	IssueDate	
4701102	Title	5301102	Sonora	Title	110587	
	Title	5001102	Sony	Title	110588	
4601102	Title	4801102	Sonora	Title	100581	
<BOM-level>		Phone Innards SubAssembly				
Source Patent	Src Pat Title	Citing Pat	Assignee	Title	IssueDate	
5003933	Title	Patent No.	Assignee	Title	110587	
	Title	Patent No.	Assignee	Title	110582	

FIG.63

66/180

Ownership Status: < >

Source Patent	Src Pat Title	Citing Pat	Assignee	Title	IssueDate
Patent Number	Title	Patent No.	Assignee	Title	DATE
	Title	Patent No.	Assignee	Title	DATE
Source Patent	Src Pat Title	Citing Pat	Assignee	Title	IssueDate
Patent Number	Title	Patent No.	Assignee	Title	DATE
	Title	Patent No.	Assignee	Title	DATE

FIG.64

67/180

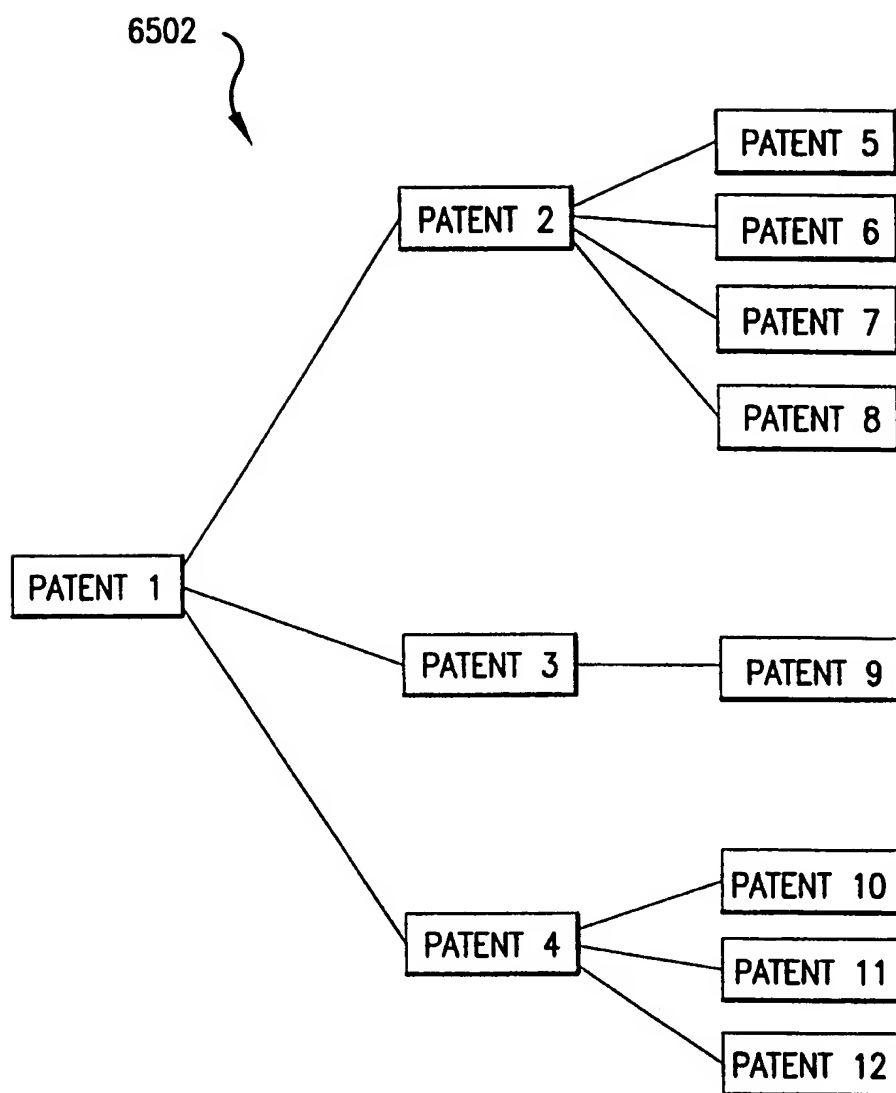
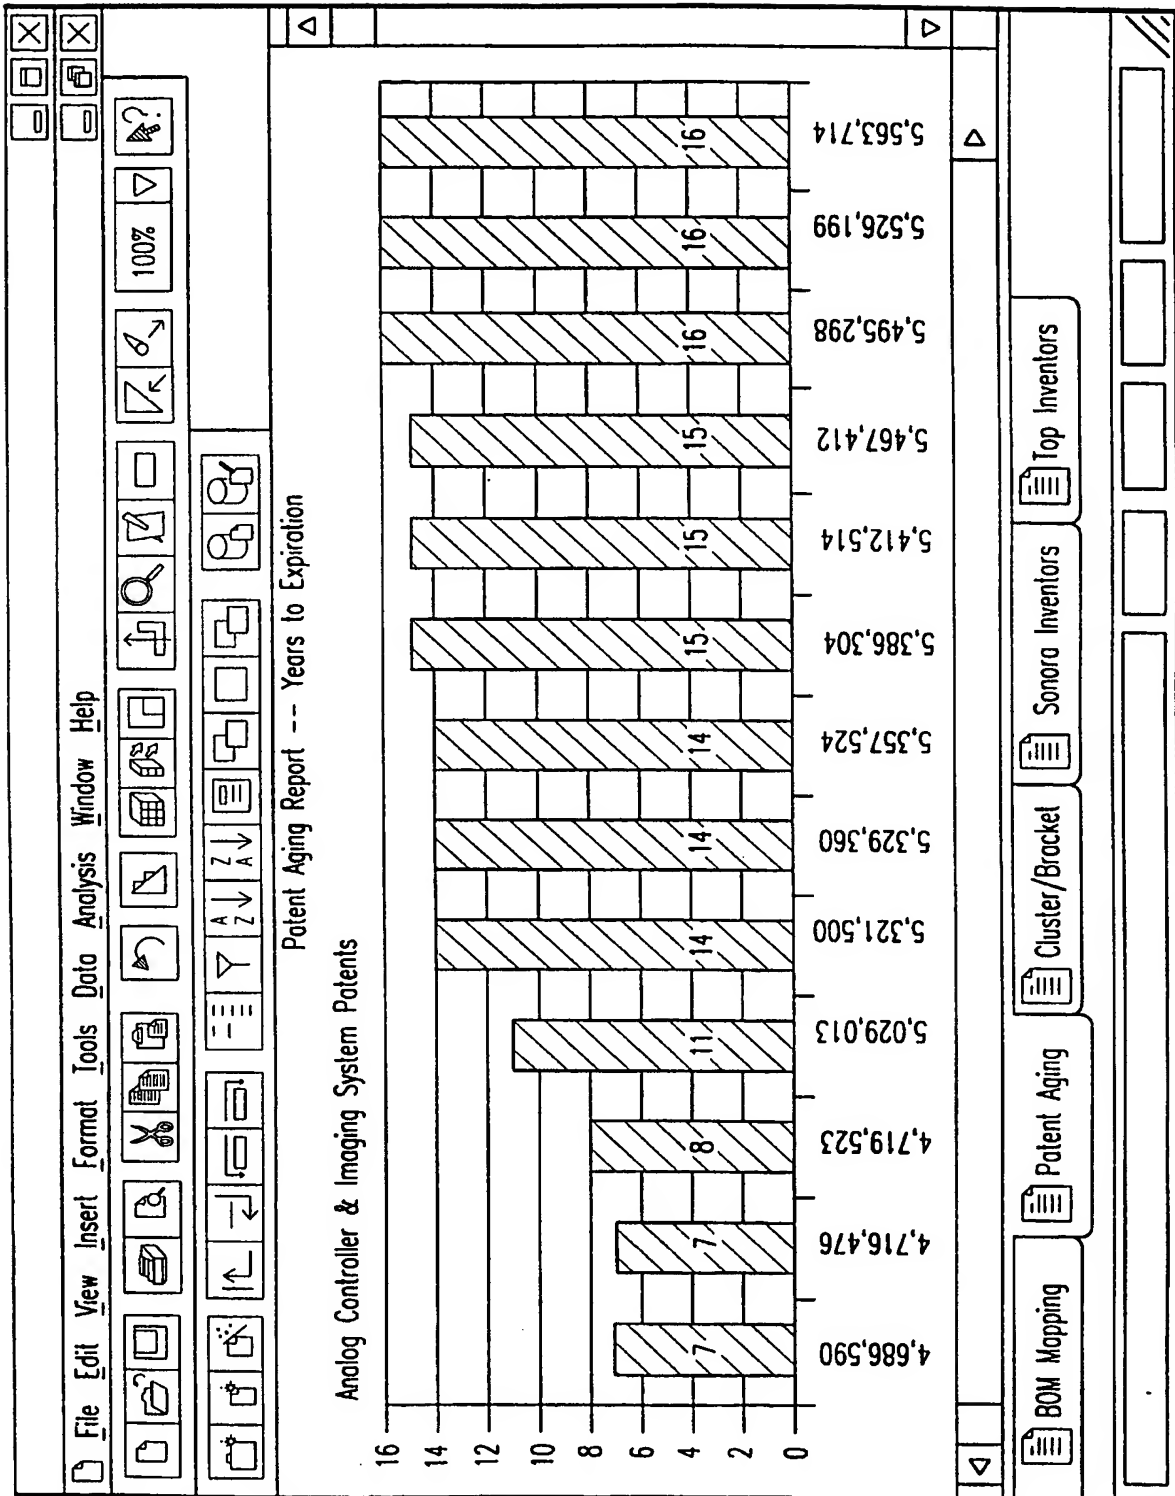


FIG.65

68/180






FIG.67

70/180

6802

<Group titles>	Patent#	Yrs to Exp
XYZ SubAssembly	5,000,123	1
	5,000,321	2
<Group titles>	Patent#	Yrs to Exp
ABC SubAssembly	5,034,123	6
	5,043,321	9
	5,073,921	9

FIG.68

6902

<Ownership Status>	Patent#	Yrs to Exp
Owned	5,000,123	1
	5,000,321	2
<Ownership Status>	Patent#	Yrs to Exp
Licensed	5,034,123	6
	5,043,321	9
	5,073,921	9

FIG.69

7002

<Org_Level Descr>		
Patent#	Yrs to Exp	Ownership
5,000,123 5,000,321	1	licensed
	2	owned
<Org_Level Descr>		
Patent#	Yrs to Exp	Ownership
5,034,123 5,043,321 5,073,921	6	owned
	9	owned
	9	licensed

FIG.70

71/180




FIG.71




FIG.72

72/180




FIG. 73

73/180

7402




FIG.74




FIG.75

<Group Title> 7602

Inventor Name	# patents	emp_status	# core patents	# non-core patents
C. Rabb		5 active		
A. Jackson		4 active		
L. Hohmann		2 inactive		

FIG.76

<Group Title> 7702

<Org Level>	Name	# patents	emp_status
Engineering	C. Rabb		5 active
	A. Jackson		4 active
	J. Smith		active
<Org Level>	Name	# patents	emp_status
Marketing	D. Puglia		2 active
	L. Hohmann		2 inactive

FIG.77

76/180

7802

Inventor Employment Information																																																					
<div>Analog Special Effects Controller SubA</div> <table border="1"> <thead> <tr> <th>Patent</th> <th>Emp ID</th> <th>LastName</th> <th>FirstName</th> <th>Department</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>5029013</td> <td>13178</td> <td>Fukushima</td> <td>Shinya</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5029013</td> <td>13202</td> <td>Hiratsuka</td> <td>Masaru</td> <td>Engineering</td> <td>Inactive</td> </tr> <tr> <td>5029013</td> <td>13309</td> <td>Nakagawa</td> <td>Kiyoshi</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5029013</td> <td>13315</td> <td>Nakashima</td> <td>Hiroshi</td> <td>Engineering</td> <td>Active</td> </tr> </tbody> </table>						Patent	Emp ID	LastName	FirstName	Department	Status	5029013	13178	Fukushima	Shinya	Engineering	Active	5029013	13202	Hiratsuka	Masaru	Engineering	Inactive	5029013	13309	Nakagawa	Kiyoshi	Engineering	Active	5029013	13315	Nakashima	Hiroshi	Engineering	Active																		
Patent	Emp ID	LastName	FirstName	Department	Status																																																
5029013	13178	Fukushima	Shinya	Engineering	Active																																																
5029013	13202	Hiratsuka	Masaru	Engineering	Inactive																																																
5029013	13309	Nakagawa	Kiyoshi	Engineering	Active																																																
5029013	13315	Nakashima	Hiroshi	Engineering	Active																																																
<div>Digital Color Correction Subsystem</div> <table border="1"> <thead> <tr> <th>Patent</th> <th>Emp ID</th> <th>LastName</th> <th>FirstName</th> <th>Department</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>5353059</td> <td>13274</td> <td>Lowlor</td> <td>Robert J.D.</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5353059</td> <td>13441</td> <td>Wilkinson</td> <td>James H.</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5381175</td> <td>13253</td> <td>Kihara</td> <td>Toku</td> <td>R&D</td> <td>Active</td> </tr> <tr> <td>5381175</td> <td>13399</td> <td>Sudo</td> <td>Fumihiko</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5452018</td> <td>13152</td> <td>Capitant</td> <td>Patrice</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5452018</td> <td>13154</td> <td>Carlucci</td> <td>John</td> <td>Engineering</td> <td>Active</td> </tr> <tr> <td>5452018</td> <td>13158</td> <td>Collier</td> <td>David</td> <td>Engineering</td> <td>Inactive</td> </tr> </tbody> </table>						Patent	Emp ID	LastName	FirstName	Department	Status	5353059	13274	Lowlor	Robert J.D.	Engineering	Active	5353059	13441	Wilkinson	James H.	Engineering	Active	5381175	13253	Kihara	Toku	R&D	Active	5381175	13399	Sudo	Fumihiko	Engineering	Active	5452018	13152	Capitant	Patrice	Engineering	Active	5452018	13154	Carlucci	John	Engineering	Active	5452018	13158	Collier	David	Engineering	Inactive
Patent	Emp ID	LastName	FirstName	Department	Status																																																
5353059	13274	Lowlor	Robert J.D.	Engineering	Active																																																
5353059	13441	Wilkinson	James H.	Engineering	Active																																																
5381175	13253	Kihara	Toku	R&D	Active																																																
5381175	13399	Sudo	Fumihiko	Engineering	Active																																																
5452018	13152	Capitant	Patrice	Engineering	Active																																																
5452018	13154	Carlucci	John	Engineering	Active																																																
5452018	13158	Collier	David	Engineering	Inactive																																																
<div> <div>BOM Mapping</div> <div>Patent Aging</div> <div>Cluster/Bracket</div> <div>Sonora Inventors</div> <div>Top Inventors</div> </div>																																																					

FIG.78

7902

77/180

<BOM_Level>					
<BOM_Level De	Patent	Empld	Name	Org_Level	Status
Imaging Device	5003922	3892	A Jackson	Engineering	Active
	5003922	1224	J Smith	Engineering	Active
	5003922	1266	D Puglia	Engineering	Active
	5103944	3892	A Jackson	Engineering	Active
	5100014	3892	A Jackson	Engineering	Active
BOM_Level					
<BOM_Level De	Patent	Empld	Name	Org_Level	Status
Cup holder	4283039	3833	C Rabb	Engineering	Active
	4739181		A Einstein		
	4983039	7132	L Hohmann	Engineering	Inactive

FIG.79

8002

Sony Electronics					
Org_Level2	Patent	Empld	Name	Status	
Engineering	5003922	3892	A Jackson	Active	
Engineering	5103944	3892	A Jackson	Active	
Sony Music					
Org_Level2	Patent	Empld	Name	Status	
	4739181		A Einstein		
R&D	4283039	3833	C Rabb	Active	

FIG.80

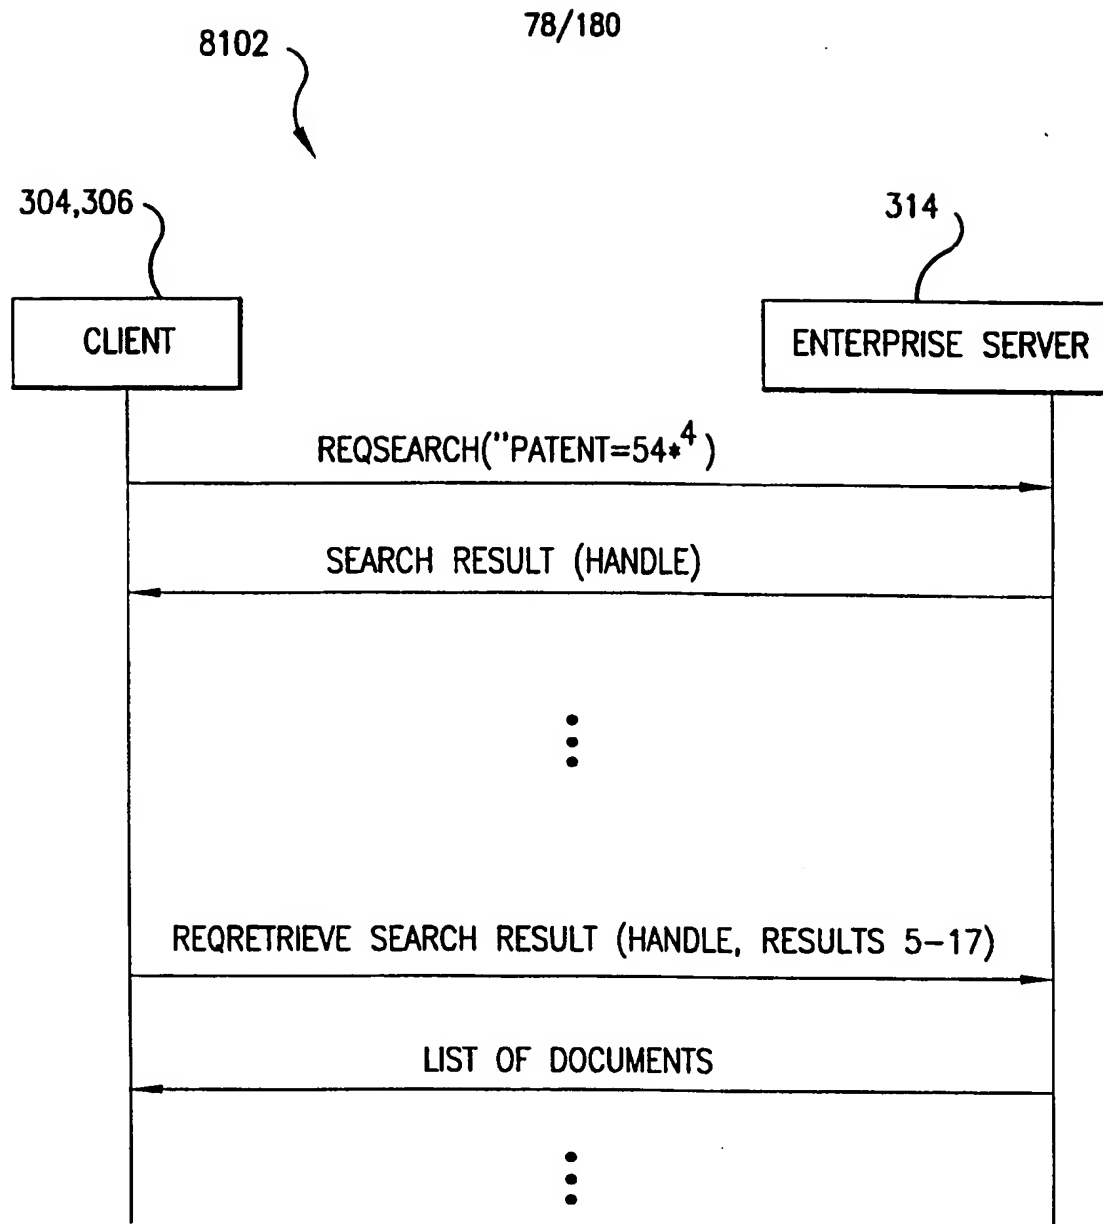


FIG.81

79/180




FIG.82

80/180




FIG.83




FIG.84

82/180






FIG.85





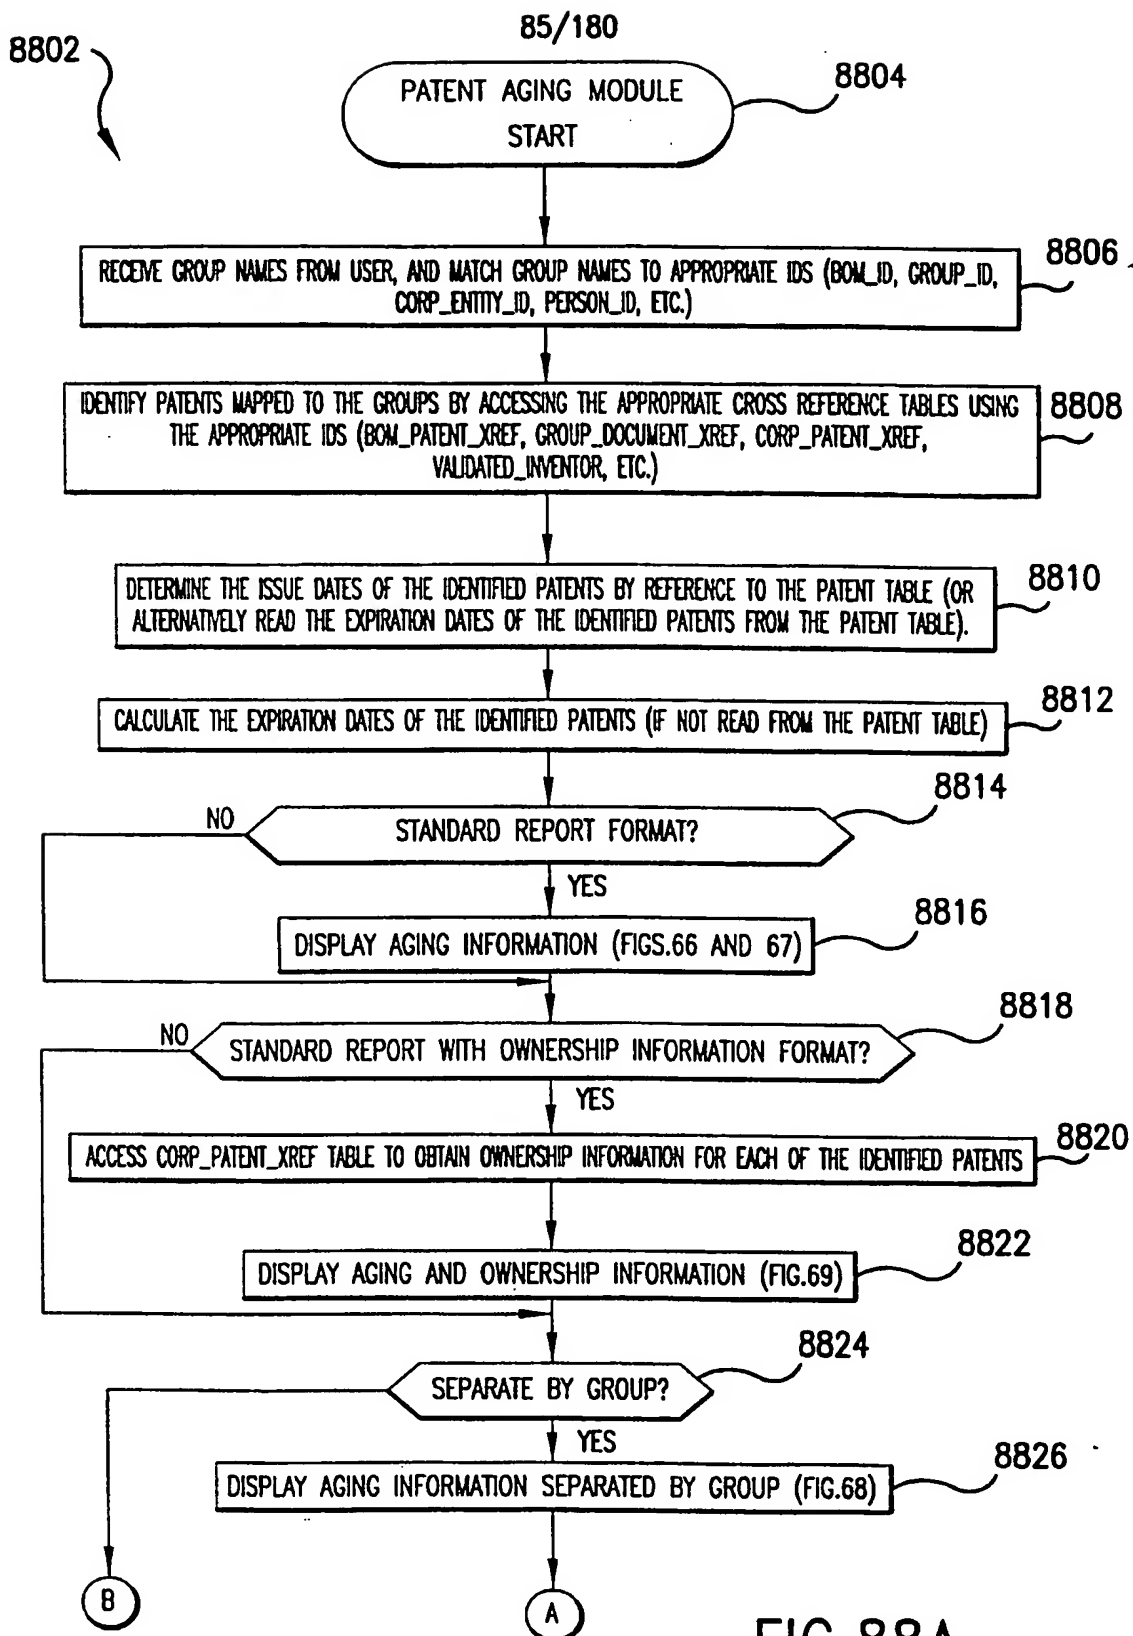


FIG.88A

86/180




FIG.88B




FIG. 89


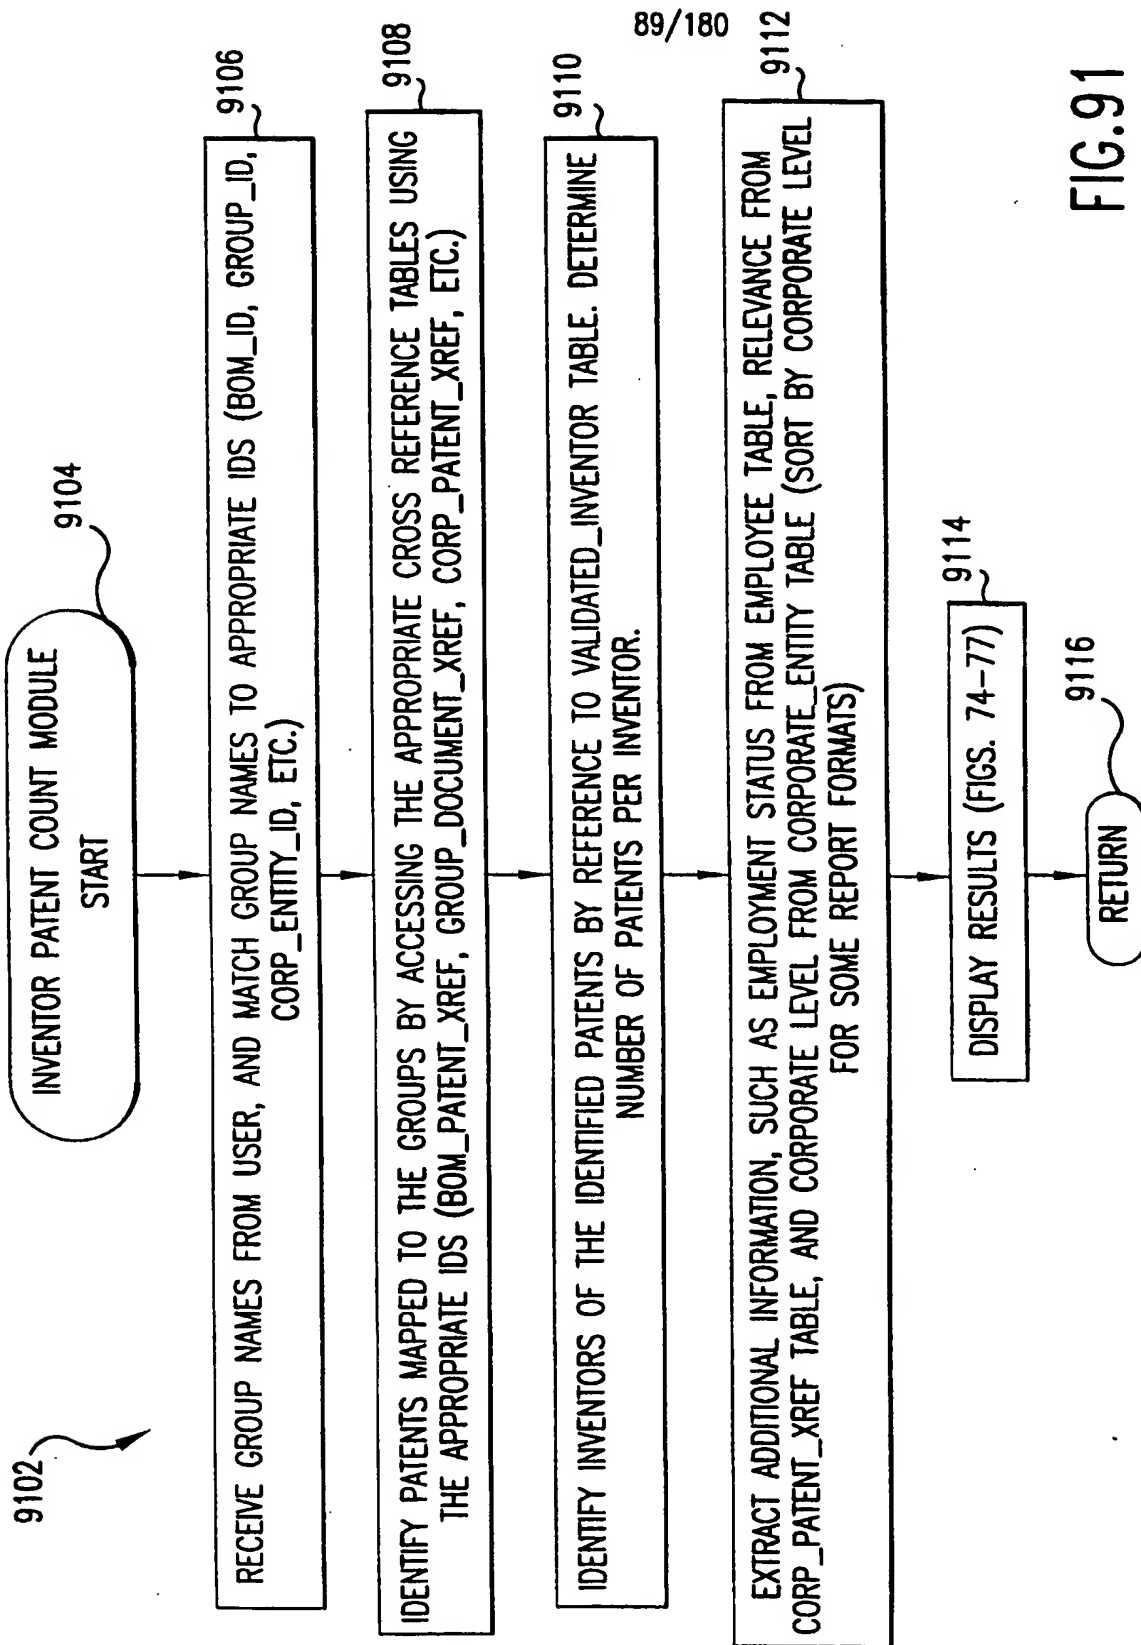



FIG. 90








FIG.93




FIG.94

93/180

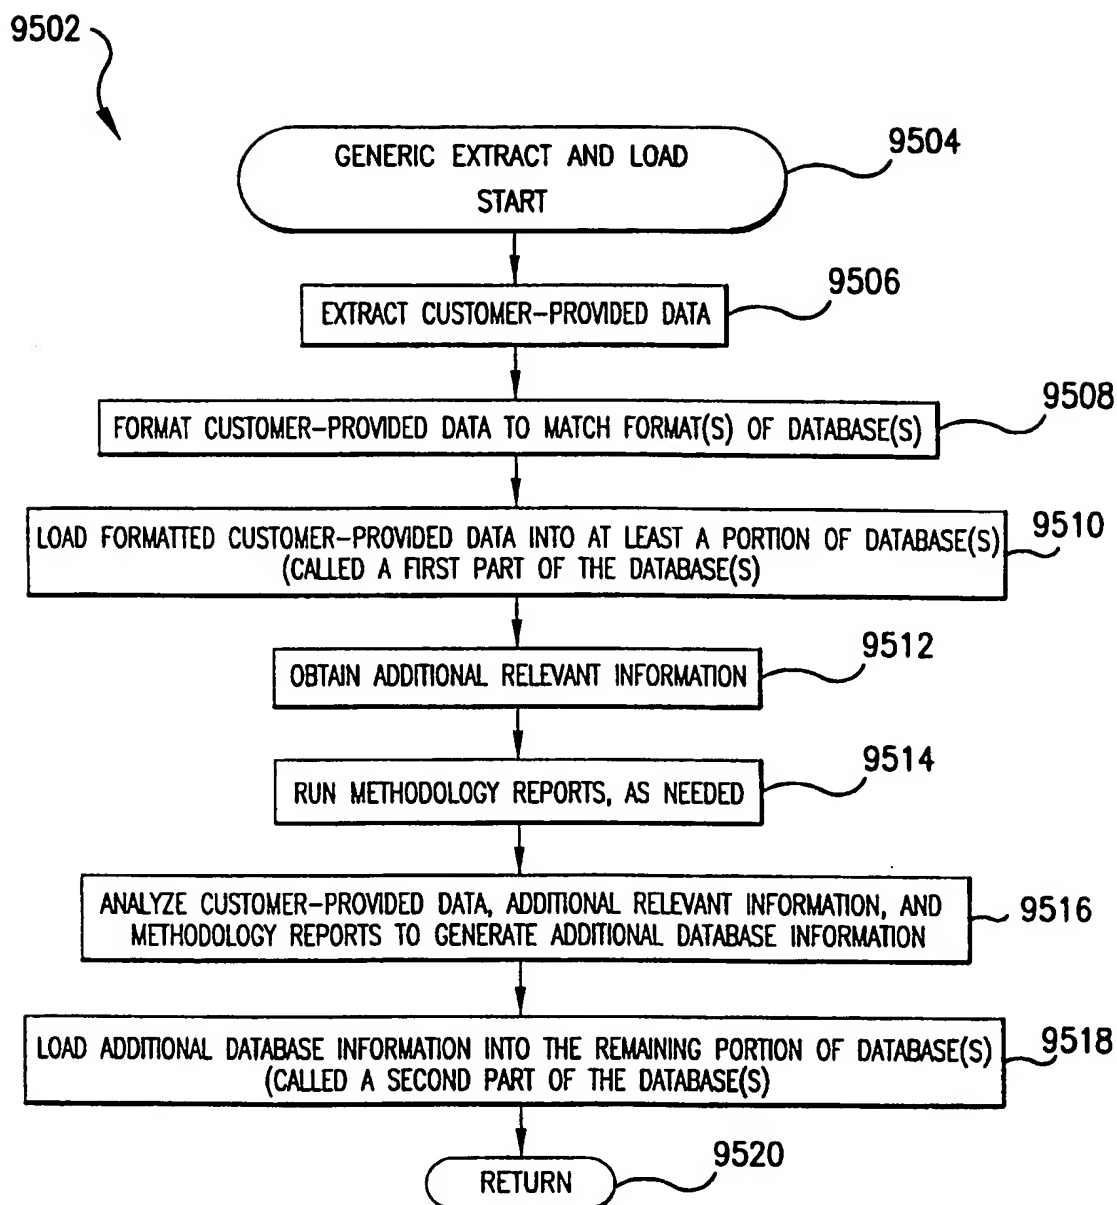


FIG.95

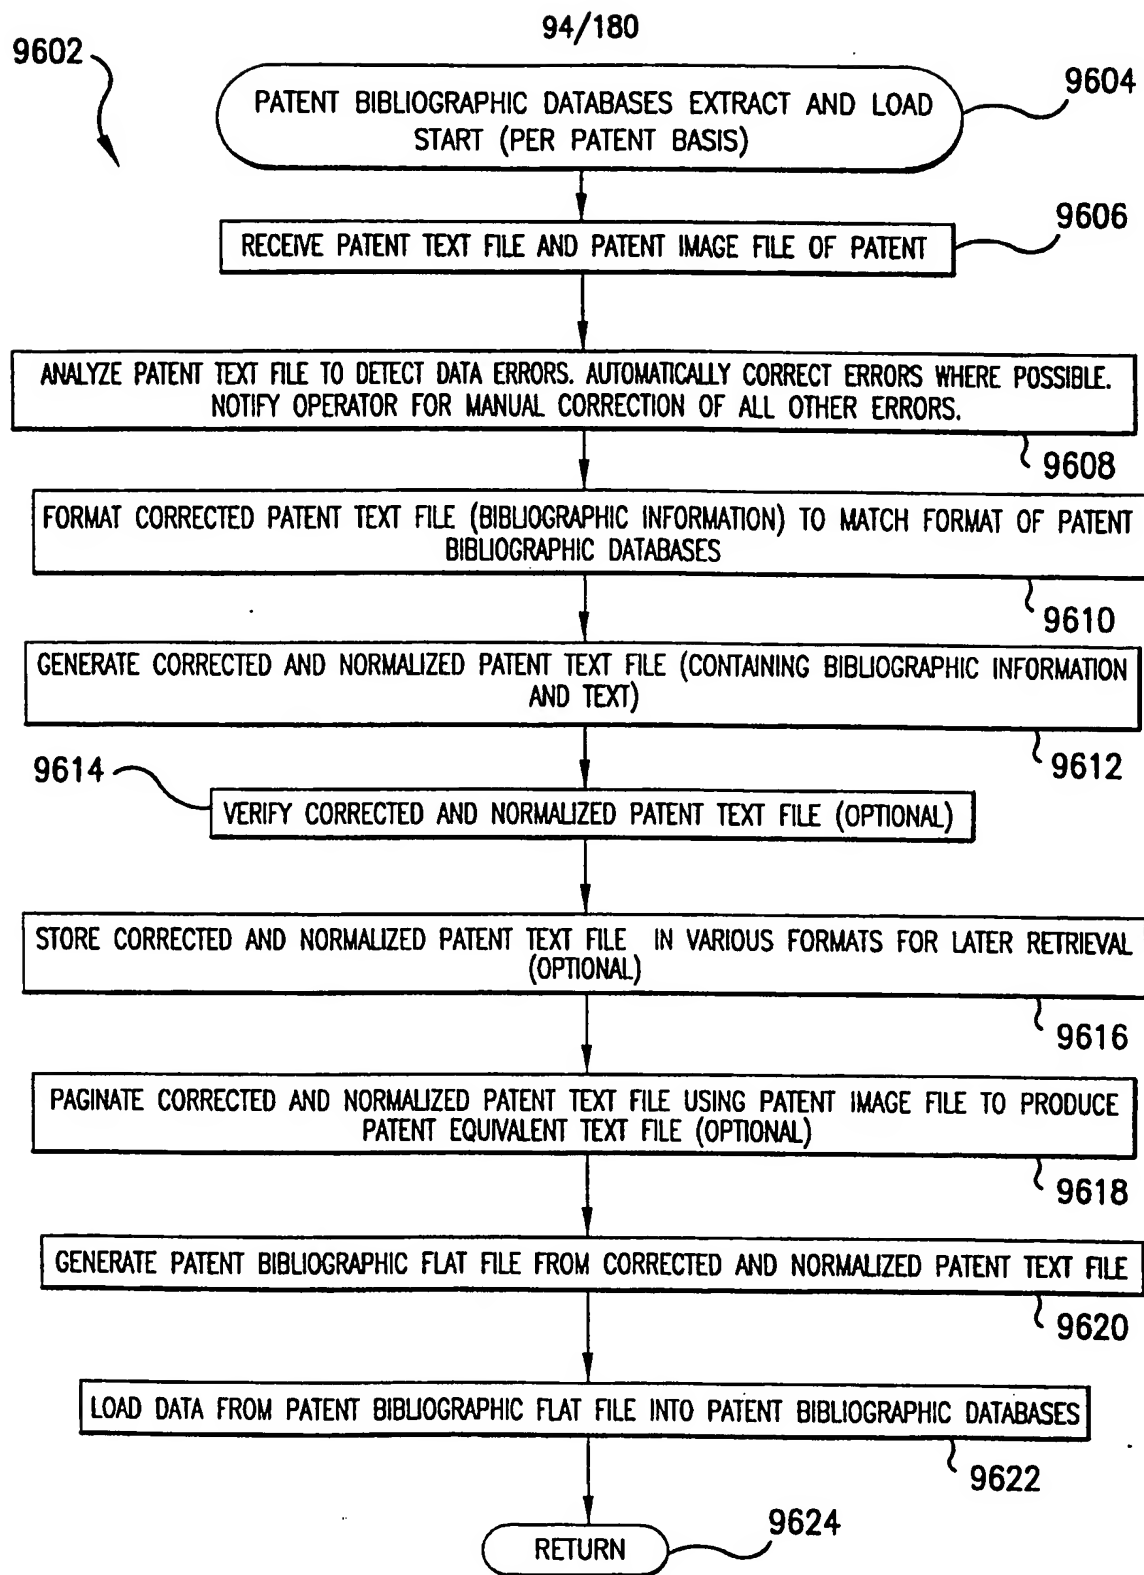


FIG.96

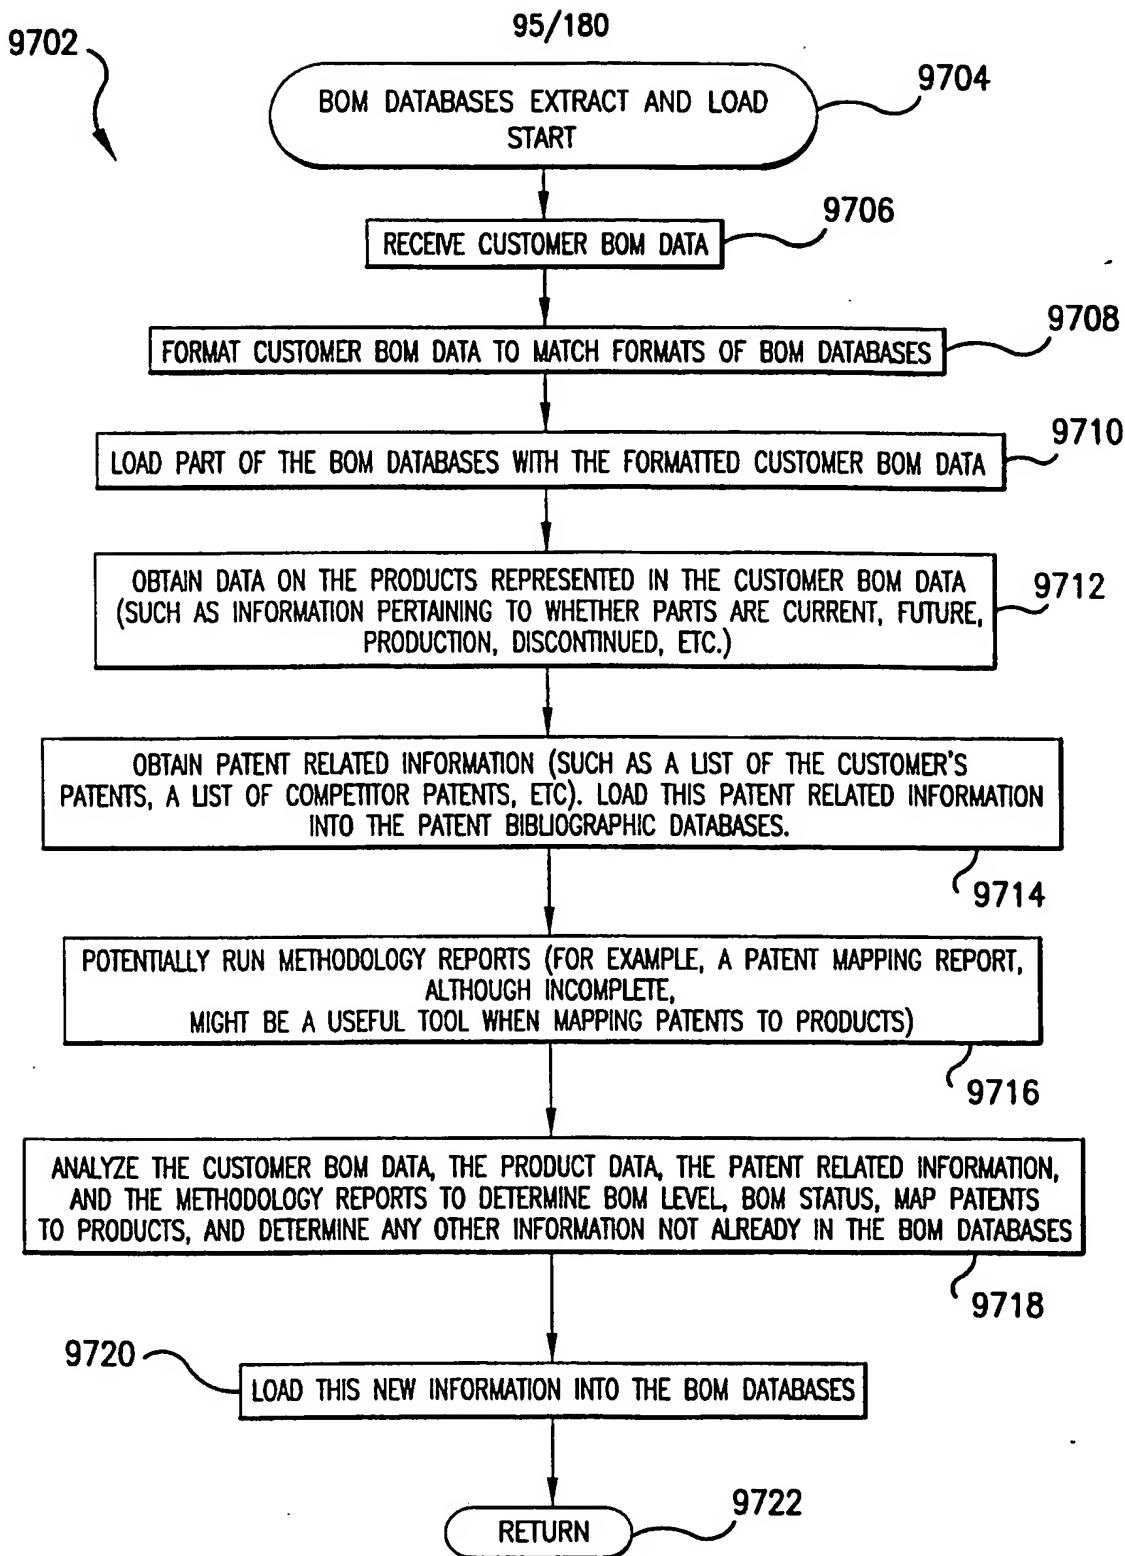
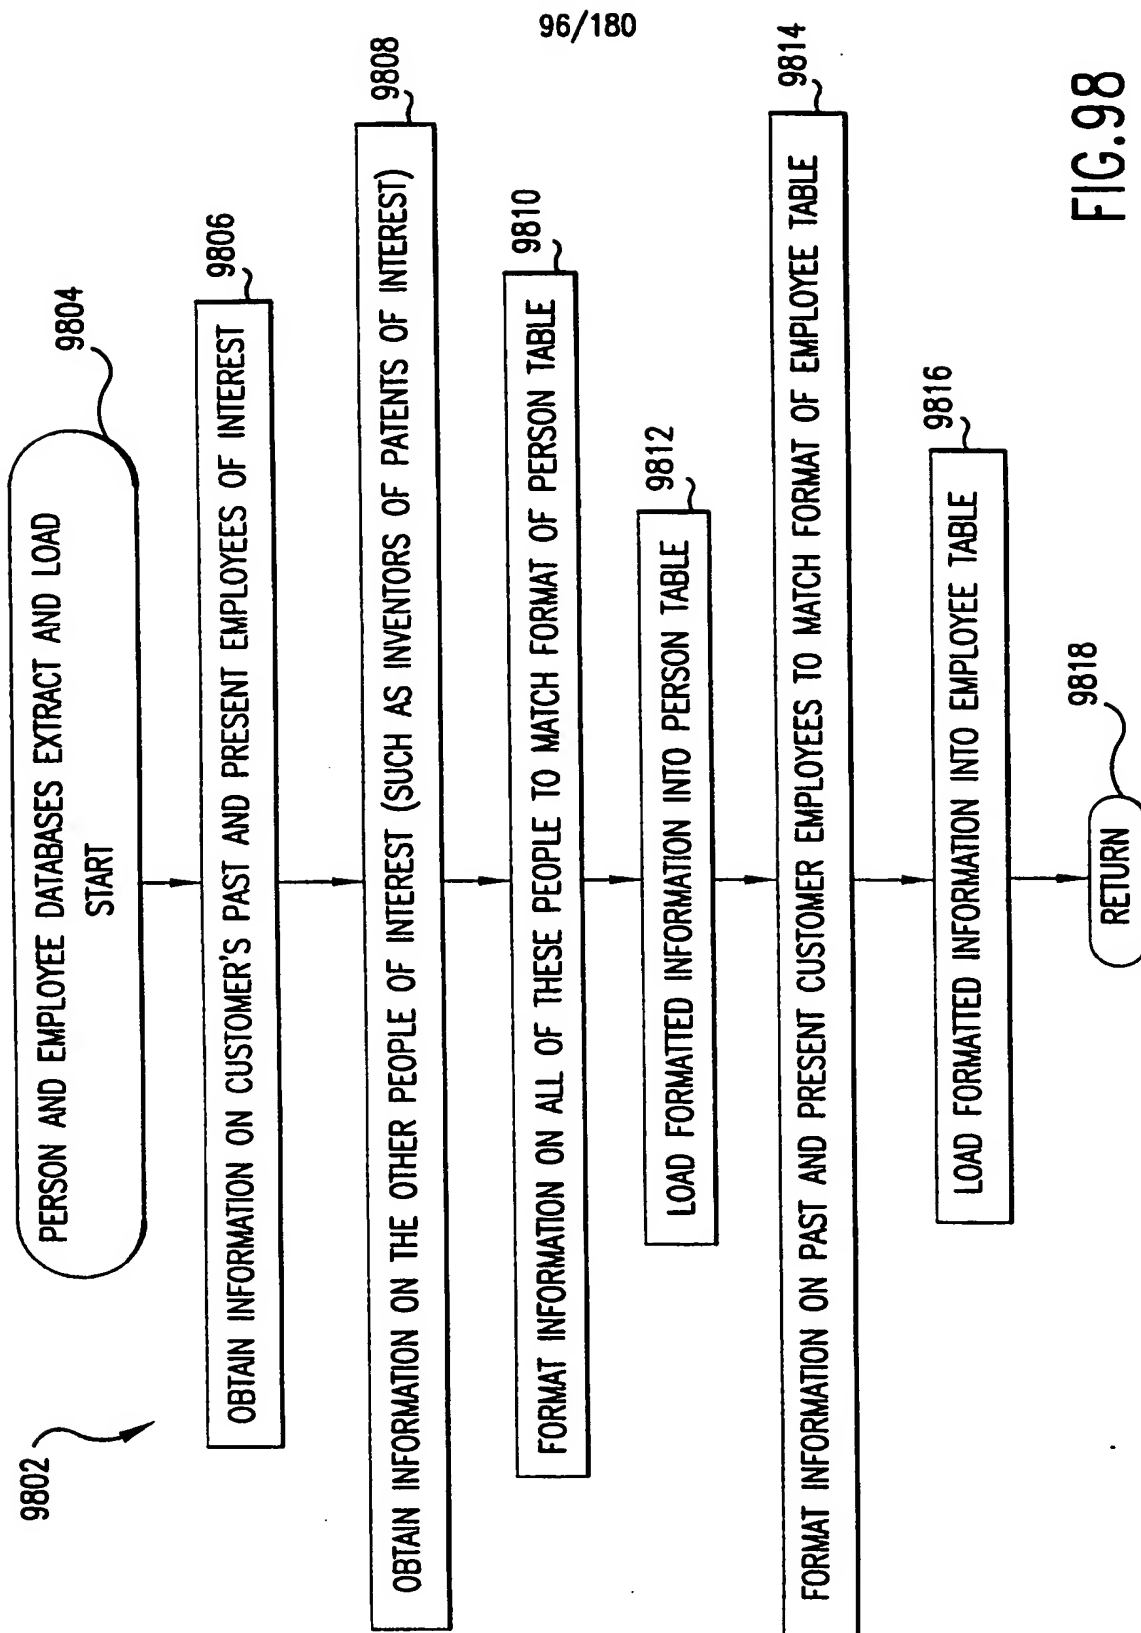


FIG.97






FIG.99

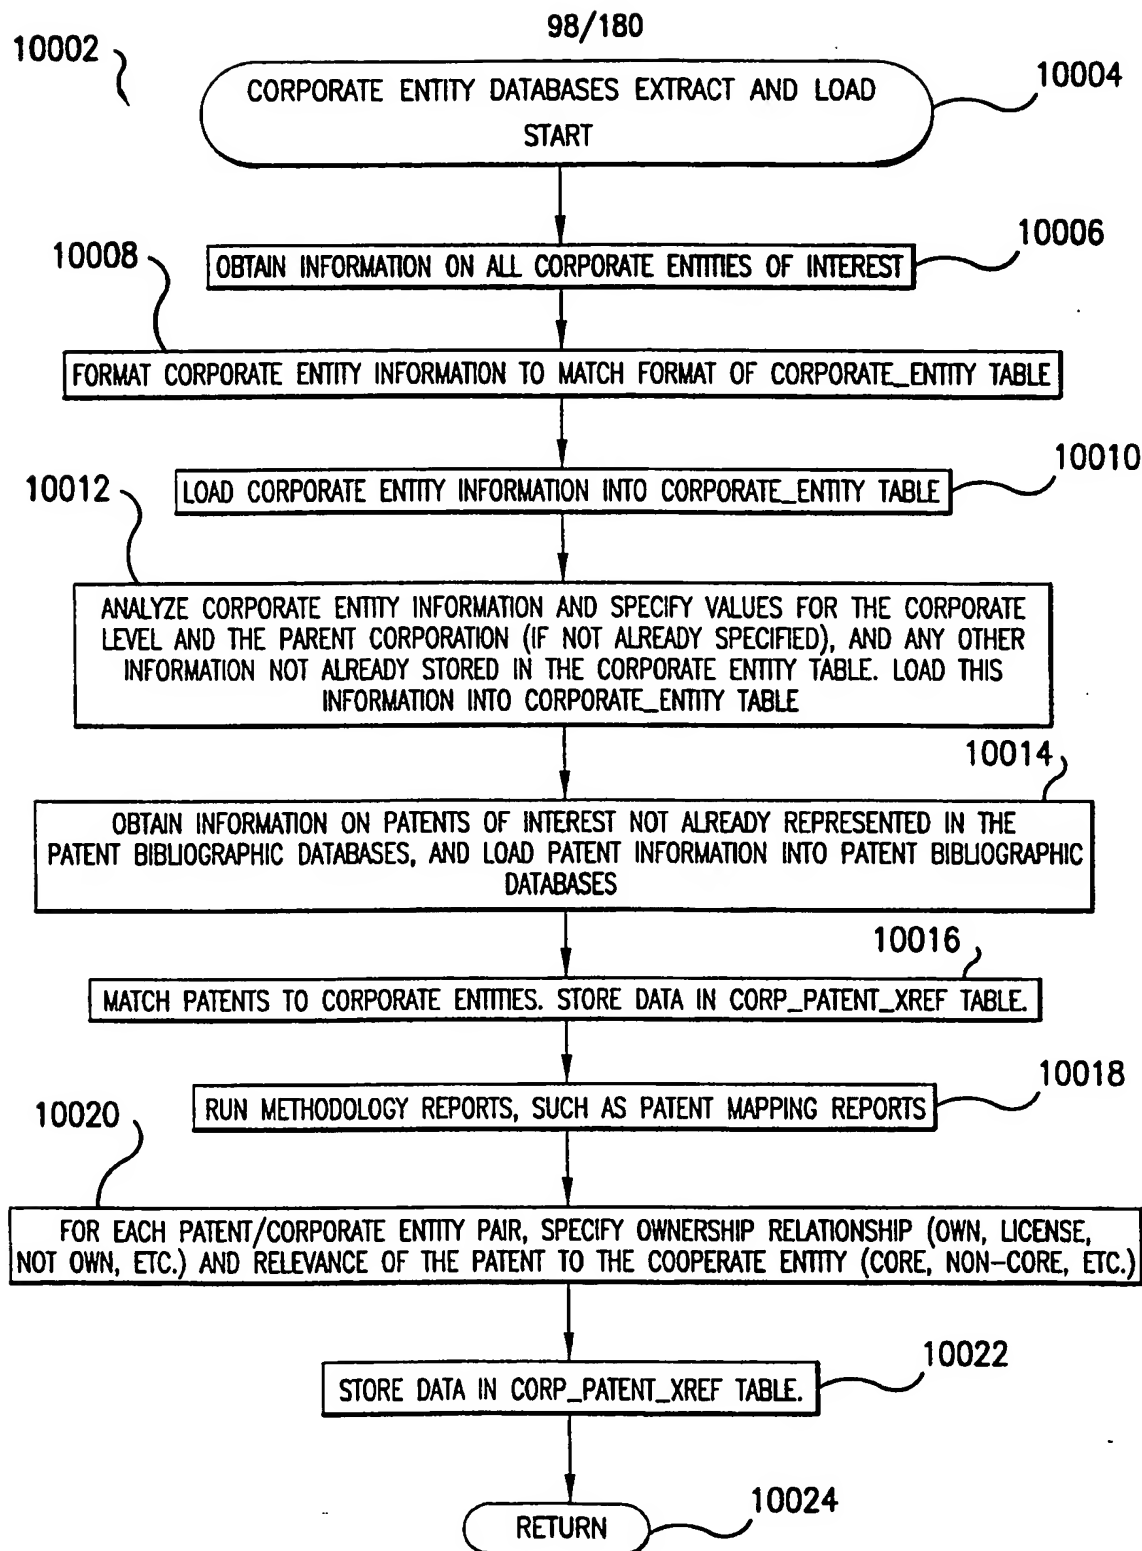


FIG.100

10102

99/180

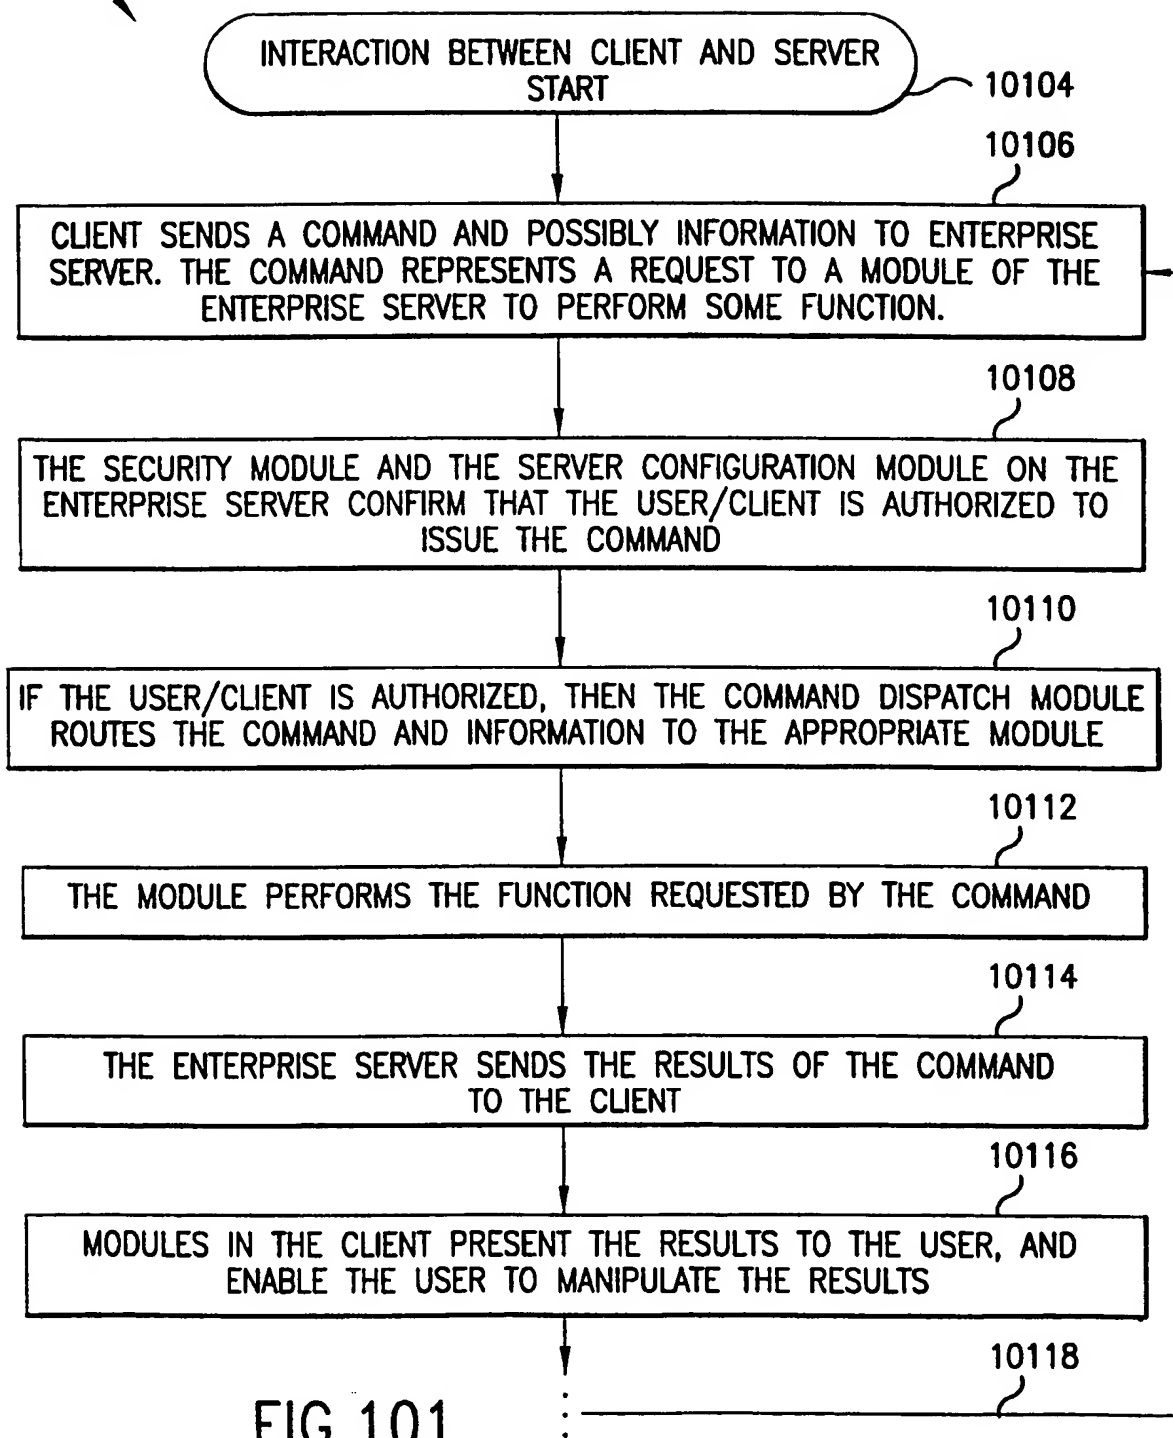
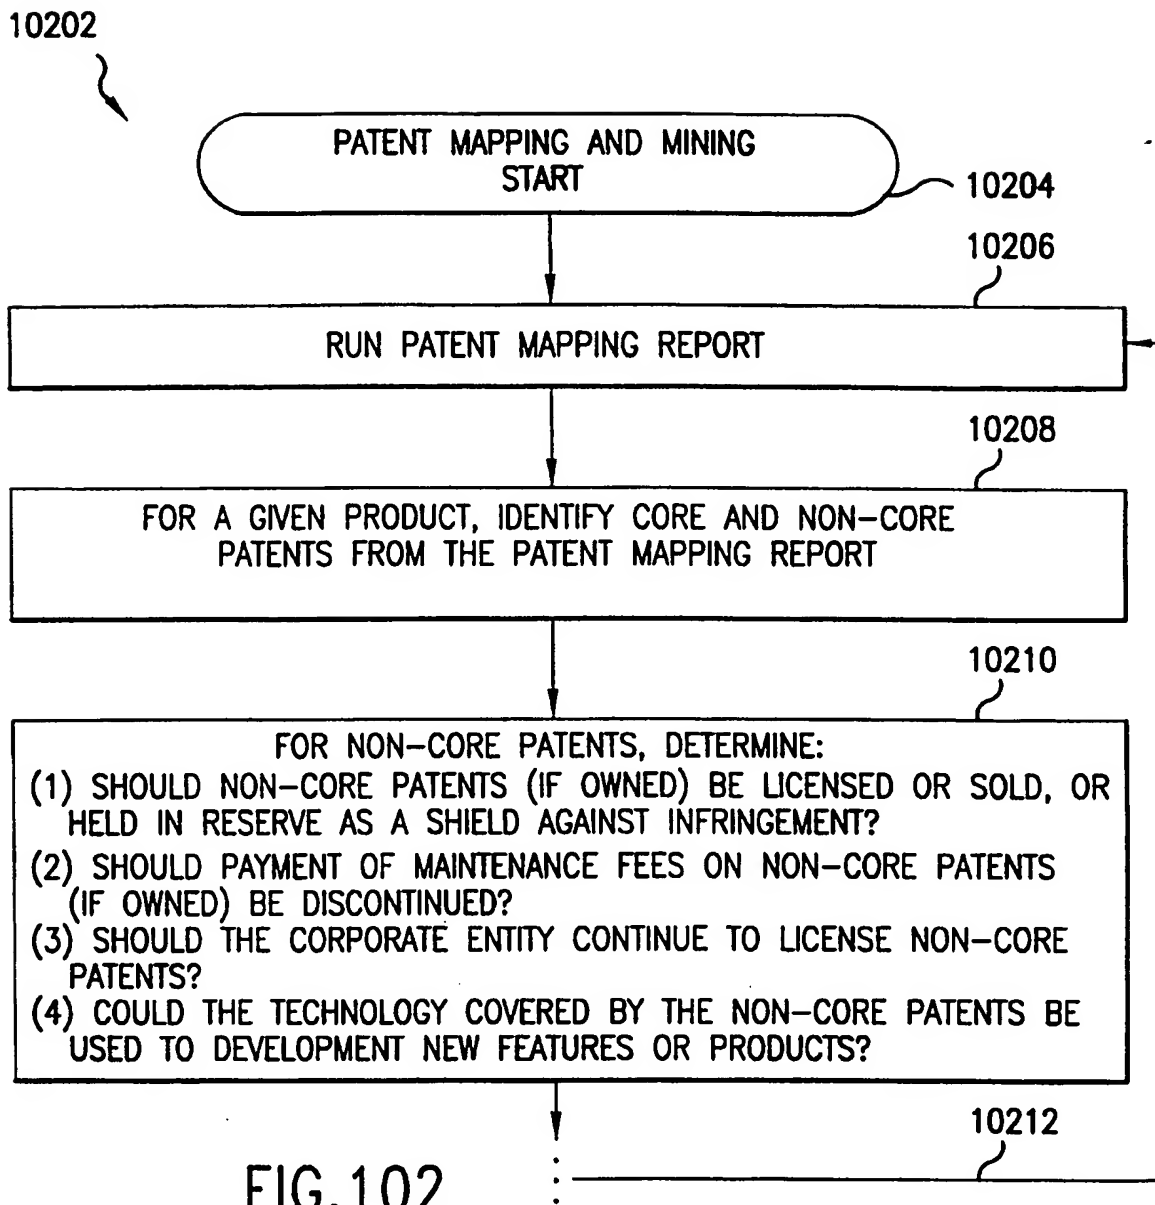


FIG.101

100/180



101/180

10302

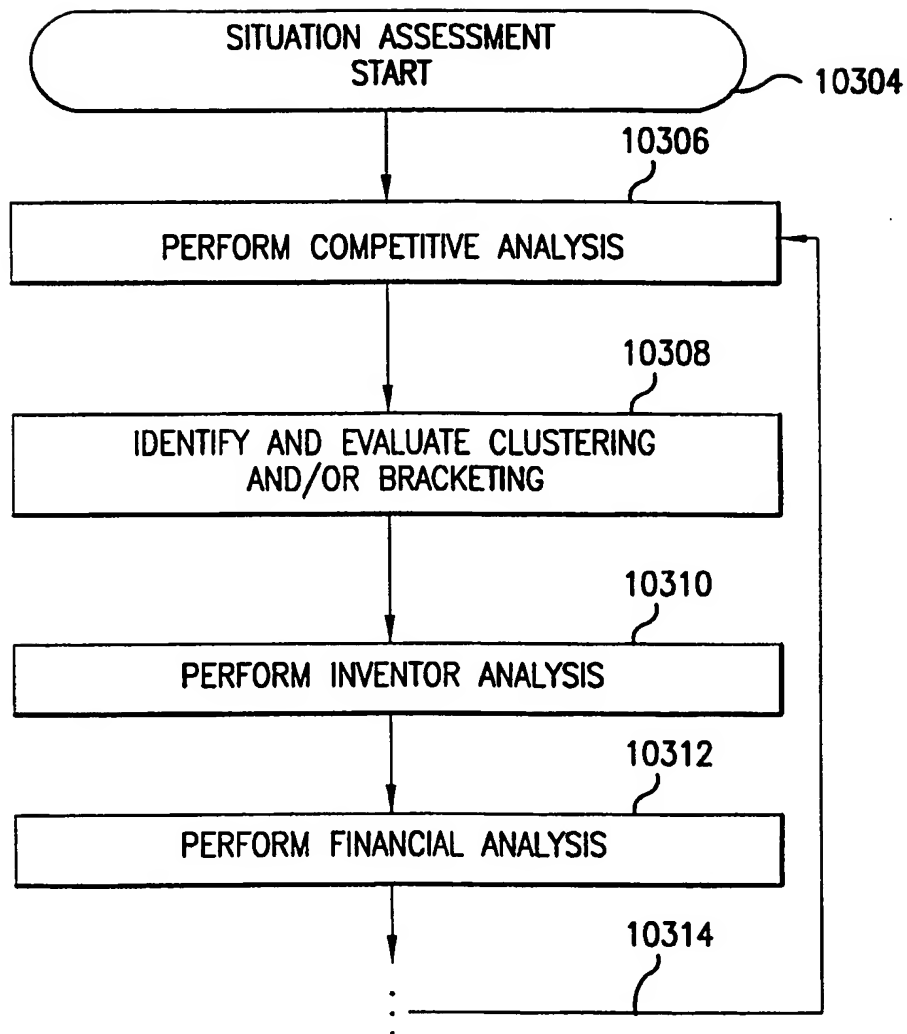
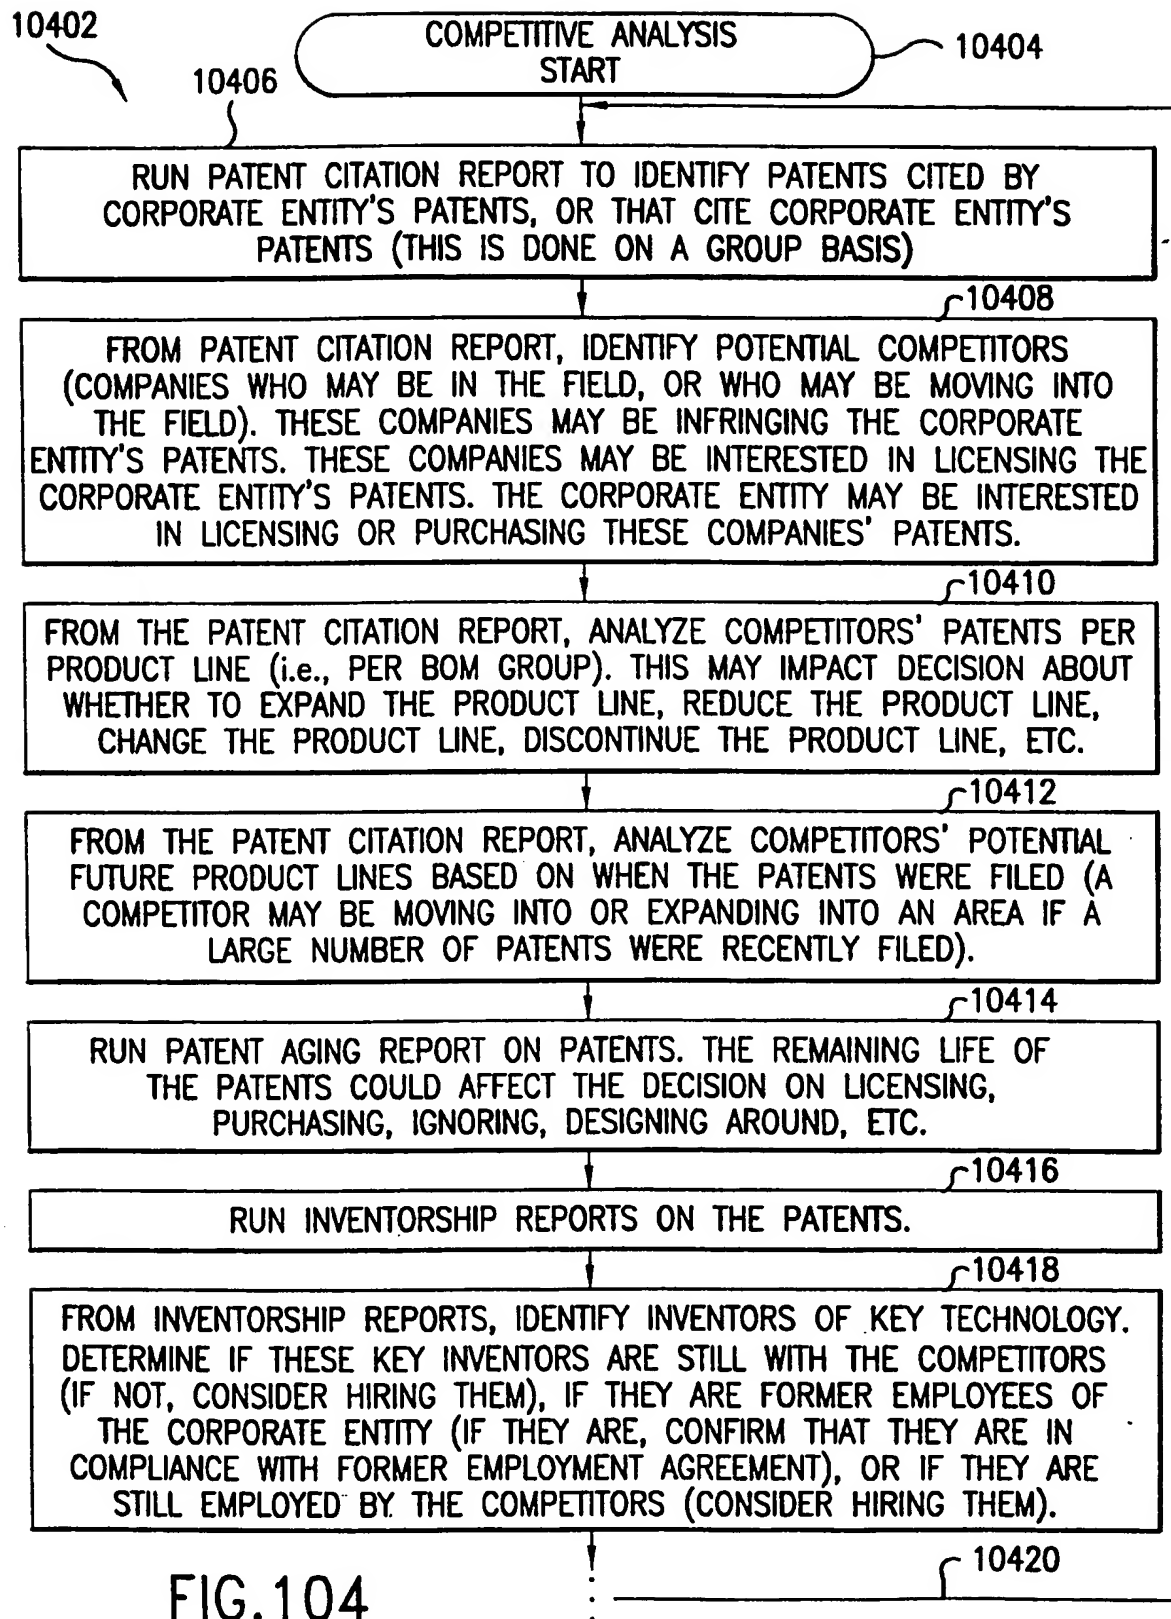


FIG.103



103/180

10502

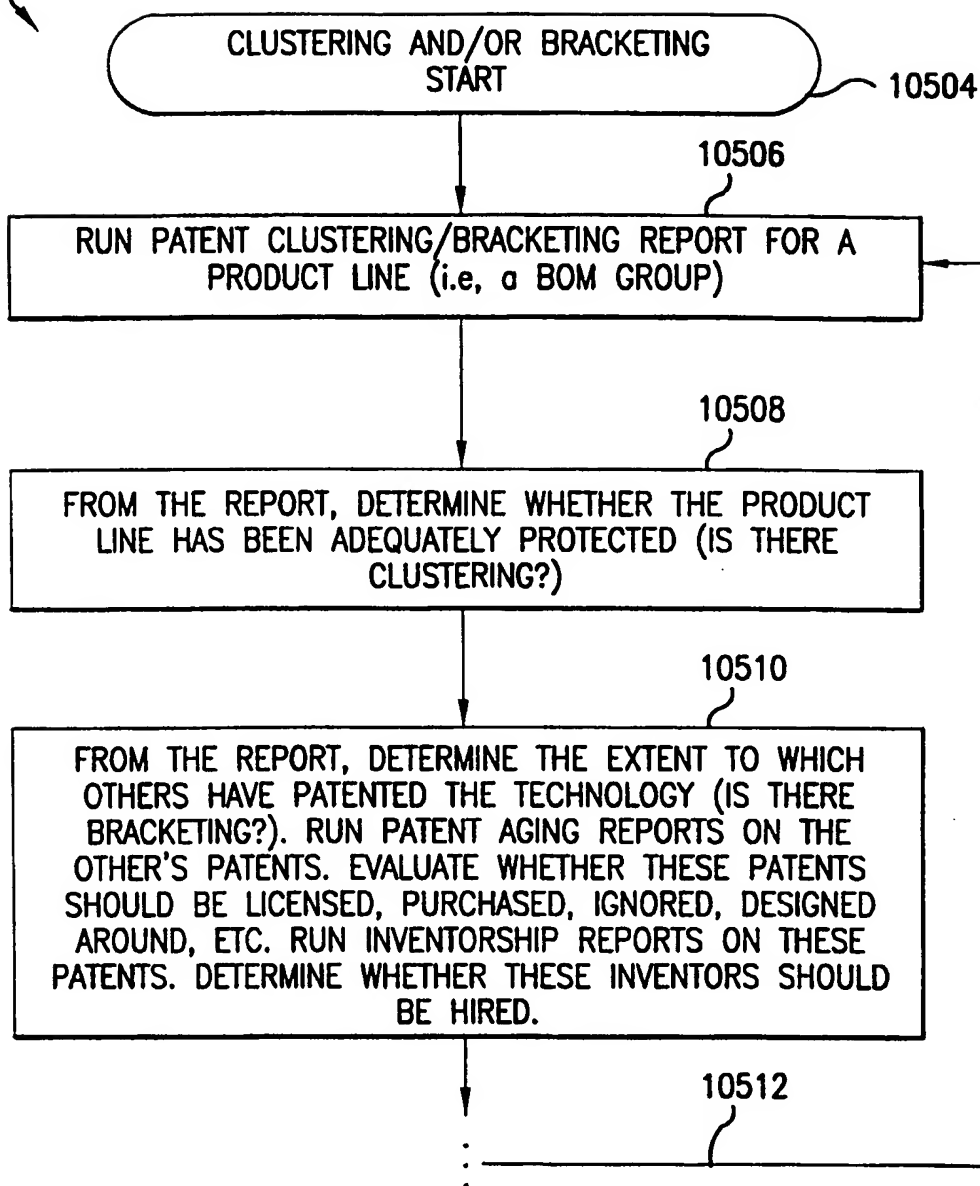


FIG.105

104/180

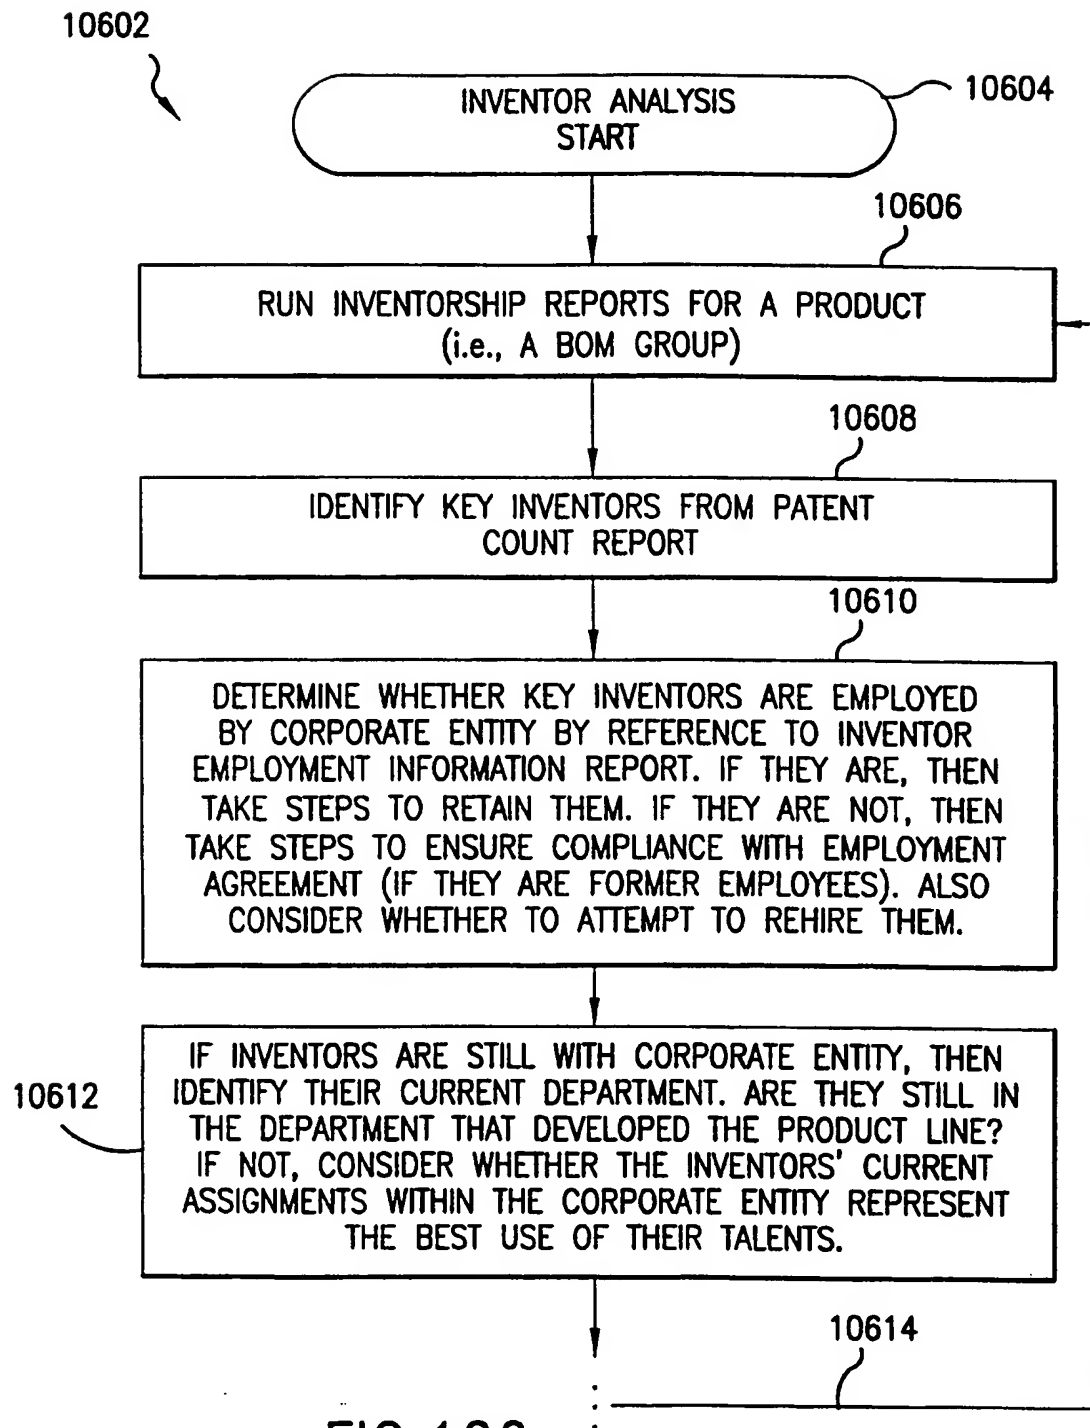


FIG.106

105/180

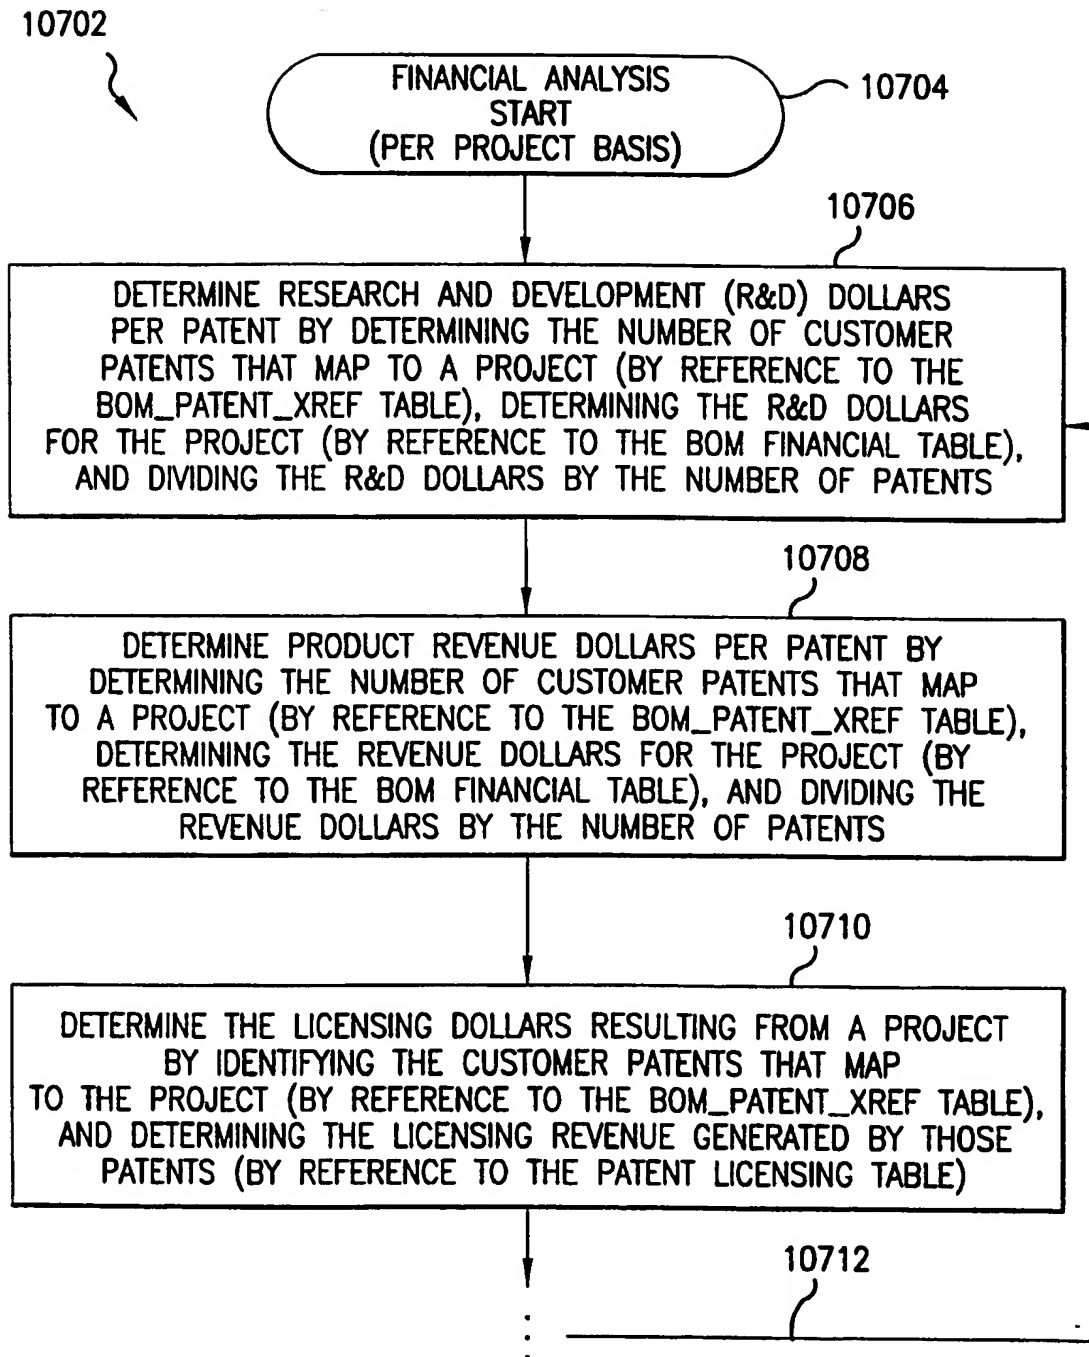


FIG.107

106/180

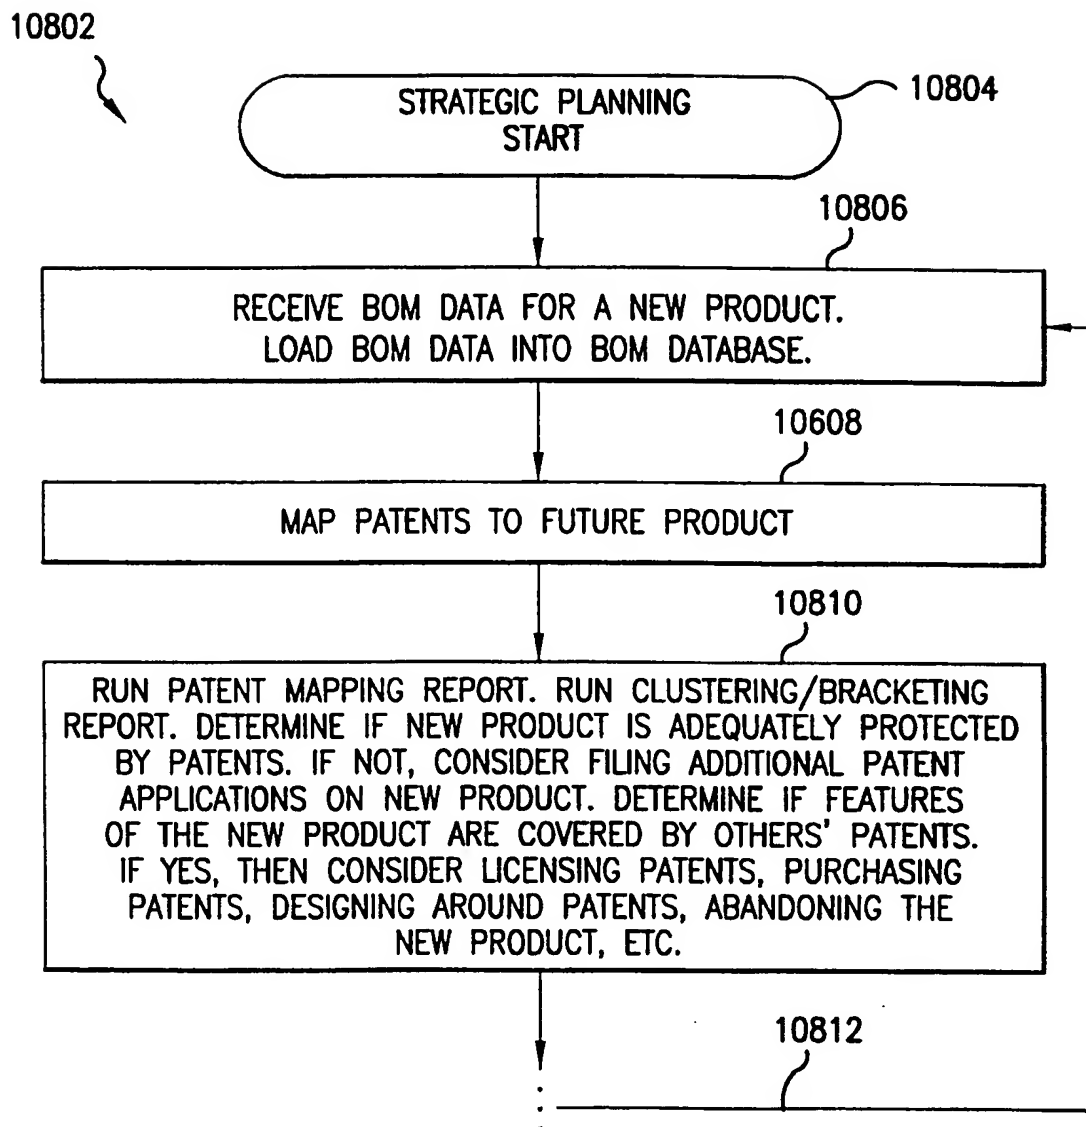


FIG.108

107/180




FIG.109

108/180

11002

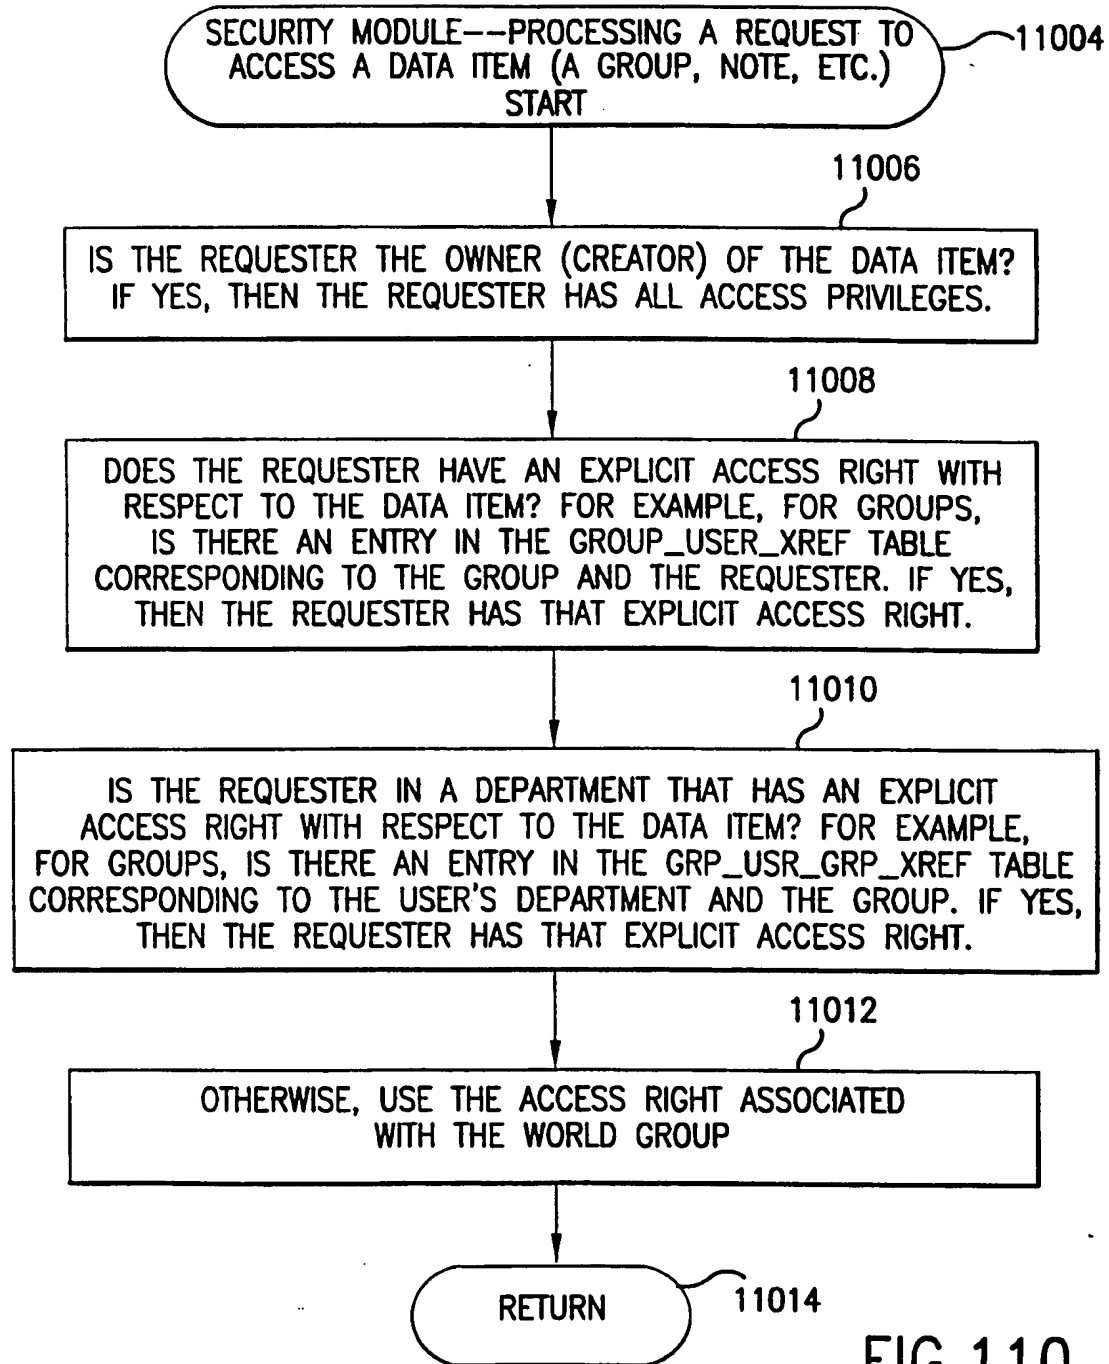


FIG.110

CASE	EDIT	PATENT	NOTE	LIBRARY	VIEW	WINDOW	HELP
<div style="display: flex; justify-content: space-between; align-items: center;"> <div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">DEMONSTRATION</div> <div style="display: flex; align-items: center;"> <div style="border: 1px solid black; padding: 2px; margin-right: 5px;">4</div> <div style="border: 1px solid black; padding: 2px; margin-right: 5px;">4,760,478</div> <div style="border: 1px solid black; padding: 2px; margin-right: 5px;">5</div> <div style="border: 1px solid black; padding: 2px; margin-right: 5px;">PATENT NOTE 3</div> <div style="border: 1px solid black; padding: 2px; margin-right: 5px;">6</div> </div> </div> <div style="text-align: center; margin-top: 10px;"> <div style="border: 1px solid black; padding: 5px; display: inline-block;">11108</div> </div> </div>							
<div style="display: flex; justify-content: space-between; align-items: center;"> <div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">4,760,478</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">COLUMN: B1 OF 6</div> </div> <div style="text-align: center; margin-top: 10px;"> <div style="border: 1px solid black; padding: 5px; display: inline-block;">11104</div> </div> </div>							
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>United States Patent [19] Pol et al.</p> <p>[11] Patent Number 4,760,478 [45] Date of Patent: Jul. 26, 1988</p> <p>[54] VISCO-ELASTICALLY DAMPED MAGNETIC HEAD SUSPENSION ASSEMBLY</p> <p>[75] Inventors: George S. Pol, Sorolago, CA; Charles G. Snyder, San Jose, CA</p> <p>[73] Assignee: International Business Machines, Armonk, NY</p> <p>[21] Appl. No.: 883,201</p> <p>[22] Filed: Jul. 8, 1986</p> <p>[63] Continuation of Ser. No. 480,997 Mar. 30, 1983, abandoned.</p> <p>[51] Int. Cl. G11B 5/48; G11B21/16</p> <p>[52] U.S. Cl. 360/104; 360/97</p> <p>[58] Field of Search: 360/102-106, 97-99, 75, 78</p> <p>[56] References Cited</p> <p>U.S. PATENT DOCUMENTS 3,725,884 4/1973. Gartein.....340/174</p> </div> <div style="width: 5%; text-align: center;">B</div> <div style="width: 45%;"> <p>[11] Patent Number 4,760,478 [45] Date of Patent: Jul. 26, 1988</p> <p>[54] VISCO-ELASTICALLY DAMPED MAGNETIC HEAD SUSPENSION ASSEMBLY</p> <p>[75] Inventors: George S. Pol, Sorolago, CA; Charles G. Snyder, San Jose, CA</p> <p>[73] Assignee: International Business Machines, Armonk, NY</p> <p>[21] Appl. No.: 883,201</p> <p>[22] Filed: Jul. 8, 1986</p> <p>[63] Continuation of Ser. No. 480,997 Mar. 30, 1983, abandoned.</p> <p>[51] Int. Cl. G11B 5/48; G11B21/16</p> <p>[52] U.S. Cl. 360/104; 360/97</p> <p>[58] Field of Search: 360/102-106, 97-99, 75, 78</p> <p>[56] References Cited</p> <p>U.S. PATENT DOCUMENTS 3,725,884 4/1973. Gartein.....340/174</p> </div> </div>							
<div style="display: flex; justify-content: space-between; align-items: center;"> <div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">4,760,478</div> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;">COLUMN: B1 OF 6</div> </div> <div style="text-align: center; margin-top: 10px;"> <div style="border: 1px solid black; padding: 5px; display: inline-block;">11102</div> </div> </div>							
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>United States Patent [19] Pol et al.</p> <p>[11] Patent Number 4,760,478 [45] Date of Patent: Jul. 26, 1988</p> <p>[54] VISCO-ELASTICALLY DAMPED MAGNETIC HEAD SUSPENSION ASSEMBLY</p> <p>[75] Inventors: George S. Pol, Sorolago, CA; Charles G. Snyder, San Jose, CA</p> <p>[73] Assignee: International Business Machines, Armonk, NY</p> <p>[21] Appl. No.: 883,201</p> <p>[22] Filed: Jul. 8, 1986</p> <p>[63] Continuation of Ser. No. 480,997 Mar. 30, 1983, abandoned.</p> <p>[51] Int. Cl. G11B 5/48; G11B21/16</p> <p>[52] U.S. Cl. 360/104; 360/97</p> <p>[58] Field of Search: 360/102-106, 97-99, 75, 78</p> <p>[56] References Cited</p> <p>U.S. PATENT DOCUMENTS 3,725,884 4/1973. Gartein.....340/174</p> </div> <div style="width: 5%; text-align: center;">B</div> <div style="width: 45%;"> <p>[11] Patent Number 4,760,478 [45] Date of Patent: Jul. 26, 1988</p> <p>[54] VISCO-ELASTICALLY DAMPED MAGNETIC HEAD SUSPENSION ASSEMBLY</p> <p>[75] Inventors: George S. Pol, Sorolago, CA; Charles G. Snyder, San Jose, CA</p> <p>[73] Assignee: International Business Machines, Armonk, NY</p> <p>[21] Appl. No.: 883,201</p> <p>[22] Filed: Jul. 8, 1986</p> <p>[63] Continuation of Ser. No. 480,997 Mar. 30, 1983, abandoned.</p> <p>[51] Int. Cl. G11B 5/48; G11B21/16</p> <p>[52] U.S. Cl. 360/104; 360/97</p> <p>[58] Field of Search: 360/102-106, 97-99, 75, 78</p> <p>[56] References Cited</p> <p>U.S. PATENT DOCUMENTS 3,725,884 4/1973. Gartein.....340/174</p> </div> </div>							

110 / 180

DEMONSTRATION

4,760,478

4,760,478

12226

11218

11220

4,760,478 COLUMN: B1 OF 6

IMAGE

11210

11212

11214

11228

United States Patent [19]

Pat et al.

[11] Patent Number 4,760,478

[45] Date of Patent: Jul. 26, 1988

[54] VISCO-ELASTICALLY DAMPED MAGNETIC HEAD SUSPENSION ASSEMBLY

[75] Inventors: George S. Pal, Saratoga CA; Charles G. Snyder, San Jose, CA

[73] Assignee: International Business Machines, Armonk, NY

[21] Appl. No.: 883,201

[22] Filed: Jul. 8, 1986

Related U.S. Application Data

Continuation of Ser. No. 480,997, Mar. 30, 1983 abandoned.

[51] Int. Cl. 4.....G11B 5/48; G11B 21/16

[52] U.S. Cl. 360/104; 360/97

Continuation of Ser. No. 480,997, Mar. 30, 1983 Abandoned.

Int. Cl.G11B 5/48;G11B 21/16 11220

U.S. Cl. 360/104;360/97

Field of Search.....360/102-106

360/97-99, 75,78

References Cited

U.S. PATENT DOCUMENTS 11222

725,884 4/1973 Carlein.....340/174

189,759 2/1980 Bauck et al.....360/104

208,684 6/1980 Janssen et al.....360/104

216,505 8/1980 Grant et al.....360/104

389,688 6/1983 Higashiyama.....360/104

399,476 7/1983 King.....360/103

408M238 10/1983 Hearn.....360/104

447,493 5/1984 Driscoll et al.....360/129




FIG. 112

111/180




FIG.1 13




FIG. 114

113/180

11502

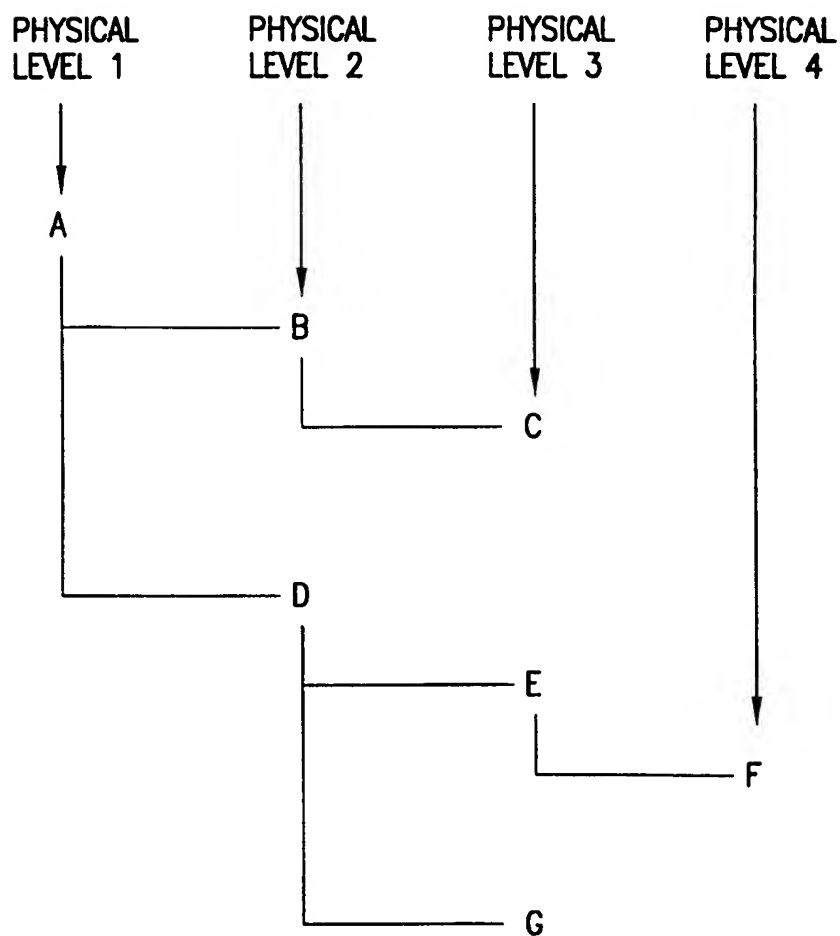


FIG.115

114/180

11602

The diagram shows a standard Windows-style dialog box titled "Login". It contains four text input fields arranged vertically. The first field is labeled "Server:" and contains the text "web.smartpatents.com:8080". The second field is labeled "Proxy:" and is empty. The third field is labeled "Login" and contains the text "su". The fourth field is labeled "Password" and contains ten asterisks. Below these fields is a checkbox labeled "Save Password" which is checked. At the bottom of the dialog are two buttons: "OK" and "Cancel". Reference numerals 11604, 11606, 11608, and 11610 point to the Server, Proxy, Login, and Password fields respectively. A reference numeral 11602 points to the entire dialog box.

Server: web.smartpatents.com:8080 11604

Proxy: 11606

Login su 11608

Password ***** 11610

☒ Save Password

OK Cancel

FIG. 116




FIG. 117

116/180

11802

web.smartpatents.com:8080

11704

11706

11720

11708

11804

11710

11712

Repository — 11710

A Group Made Remotely

A Potential Patent

Analog Video Editing Systems

Analog Professional Video Editing

Analog Special Effects Control

Video Color Correction Subsy

Video Reproduction and Copy

Digital Cameras

Disposable Cameras

Internet Appliances

Document#	Title	Assignee	Inventor	Type	Title
1 3651276	AUTOMATIC PHASING OF SERVO S	Ampex Corporation, Redwood	Harold V. Cla	1	11708
2 3682055	LENS CAP FOR CAMERA	Konan Camera Laboratory Co	Takeshi Okan	9	
3 3686432	RUNNING-SPEED CONTROL SYSTE	Matsushita Electric Industrial	Masahiro De	1	
4 3705699	TAPE RECORDER AND CASSETTE F	Grundig E.M.V., Elektro-Mech	Imre Siller, Nua	2	
5 3767853	AUTOMATIC IRIS CONTROL		Sidney Leo B	1	
6 3814854	METHOD OF SYNCHRONIZING TELE	Dalavision, Inc., Rockville, MD	Philip K. Edw	1	
7 3878560	Signal Processing and reproducing	Westinghouse Electric Corpor	William W. Ra	3	
8 3900885	Television signal time base correcto	Consolidated Video Systems,	Michael W. T	3	
9 3906166	Radio Telephone System	Motrola, Inc., Chicago, IL	Martin Coope	1	
10 3917990	Battery charging control using lemp	General Electric Company, Ly	Ralph R. She	3	
11 3925812	Color television camera including a	U.S. Philips Corporation, New	Hendrik Blom,	3	
12 3931636	Battery conservation system for col	Robert Bosch Fernsehlanlage	Hans Dieter	3	
13 3939491	Tape apparatus including means to	Sony Corporation, Tokyo, Jap	Yasuo Shigel	3	

FIG.118




FIG. 119

11704

118/180

12002

web.smartpatents.com:8080							

Document#	Title	Assignee	Inventor

Type	Title

Repository

- A Group Mode Remotely
- A Potential Patent
- Analog Video Editing Systems
 - Analog Professional Video Editing
 - Analog Special Effects Con
 - A3333 3460 M-Special
 - A3333 L193 B-Error C
 - A3333 1428 K-Special
 - A3333 8 231 Lb-High S
 - Sample New Name
 - + Video Color Correction Subs
 - + Video Reproduction and Co
 - + Digital Cameras

FIG. 120

Search			
Bib Fields	Search What?		
Patent # :	<input type="radio"/> Full Text Index <input type="radio"/> Repository <input type="radio"/> Current Group <input type="radio"/> Current Patent		
Title:			
Assignee:			
Issue Date:	<input type="text"/> <input type="button" value="v"/>	<input type="text"/> And	<input type="text"/>
Keywords:			
Class:			
Inventor:			
		<input type="button" value="Search"/>	<input type="button" value="Cancel"/>
		<input type="button" value="Save"/>	<input type="button" value="Load"/>

Text	
Abstract:	video tape <u>12104</u>
Full Text:	
<input type="button" value="AND"/> <input type="button" value="OR"/> <input type="button" value="NOT"/> <input type="button" value="w/5 chars"/> <input type="button" value="w/10 chars"/> <input type="button" value="w/25 chars"/>	

FIG. 121

120/180

12202

web.smartpatents.com:8080		TIDQ		Title		Assignee		Inventor		Type		Title	
Document#	Title	Assignee	Inventor	Type	Title								
1 3900885	Television signal time base correction	Consolidated Video Systems, Inc.	Michael W. T. 3										
2 3939491	Tape apparatus including means to	Sony Corporation, Tokyo, Jap	Yasuo Shiget 3										
3 3959818	Servo for video tape apparatus with	Sony Corporation, Tokyo, Jap	Kazuo Ikelaki 3										
4 3974522	Electronic editing method and appar	Akai Electric Company Limited	Motonori Fuk 3										
5 4001882	Magnetic tape editing, previewing a	Spectra-Vision Corporation, P	Anthony M. F 3										
6 4096523	Color correction system		Armand Belm 3										
7 4210940	Variable format tape replay system	Bell & Howell Company, Chica	Daniel G. Pry 3		11708								
8 4272780	Film scanner with variable rate imag	Teletronics International, Inc.,	Armand Belm 3										
9 4272790	Video tape editing system	Convergence Corporation, Iri	George W. B 3										
10 4308563	Control apparatus for magnetic reco	Sony Corporation, Tokyo, Jap	Seiichi Gohd 3		12208								
11 4316224	Magnetic tape reproducer-recorder	Blaupunkt-Werke GmbH, Hfde	Jens Hansen 3										
12 4366510	Video signal recording and/or repro	Sony Corporation, Tokyo, Jap	Yuji Watanab 3										
13 4394694	Tape synchronizing apparatus	Sony Corporation, Tokyo, Jap	Ichiro Ninomi 3										

12210


11706

11704

FIG. 122

121/180

12302



[XL]5,029,013 Block 1 OF 0, Line 1

United States Patent [19] [1/0]
Hiratsuka et al.

[11] Patent Number: 5,029,013
[45] Date of Patent: Jul. 2, 1991

----- ↑

[54] APPARATUS FOR EDITING VIDEO TAPES

[75] Inventors: Masaru Hiratsuka, Kanagawa, Japan; Hitoshi Nakashima, Kanagawa, Japan; Kiyoshi Nakagawa, Kanagawa, Japan; Shinya Fukushima, Saitama, Japan ↑

[73] Assignee: Sony Corporation, Tokyo, Japan ↑

[21] Appl.No.: 275,377 ↑

[22] Filed: Nov. 23, 1988

Related U.S. Application Data

[30] Foreign Application Priority Data ↑
Nov. 26, 1987 Japan.....62-298604

FIG.123

122/180

<p>5,029,013 (Page 1 of 18)</p>									
<p>United States Patent [19] Hiratsuka et al.</p>									
<p>[11] Patent Number: 5,029,013 [45] Date of Patent: Jul. 2, 1991</p>									
<p>FOREIGN PATENT DOCUMENTS</p>									
<p>[54] APPARATUS FOR EDITING VIDEO TAPES</p>									
<p>[75] Inventors: Masaru Hiratsuka ; Hitoshi Nakashima; Kiyoshi Nakagawa, all of Kanagawa; Shinya Fukushima, Saitama, all of Japan</p>									
<p>[73] Assignee: Sony Corporation, Tokyo, Japan</p>									
<p>[21] Appl.No.: 275,377</p>									
<p>[22] Filed: Nov. 23, 1988</p>									
<p>[30] Foreign Application Priority Data</p>									
<p>Nov. 26, 1987 [JP] Japan.....62-298604</p>									
<p>[51] Int. Cl.2.....E104N 5/7H2;G11B 5/86</p>									
<p>[52] U.S. Cl.358/355; 360/13; 360/15; 360/23.1</p>									
<p>[53] Field of Search358/355,311,183,22; 360/14.1, 53.1 32,15</p>									
<p>[57] ABSTRACT</p>									
<p>Dual video recording/reproducing apparatus including a video playback unit, a video recording unit having recording circuitry for recording a video output signal from the video playback unit, as well reproducing circuitry for reproducing a video signal, and a sub-picture signal former responsive in a video signal supplied thereto for providing a sub-picture insert of a main picture. A switch couples either the video output from the video playback unit or the video output from the reproducing circuitry included in the video recording unit in the sub-picture signal former. The video signal which is not</p>									

FIG.124

12402

12502 → 11706

web.smartpatents.com:8080

Repository

- A Group Made Remotely
- A Potential Patent
- Analog Video Editing Systems

[X] 5,029,013 Block 1 OF 0, Line 1

Document# Title Assignee Inventor

1	5029013	Apparatus for editing video tapes	Sony Corpora	Masaru Hirat
---	---------	-----------------------------------	--------------	--------------

12504

Simple Note

Simple Note

Ally. Work Prod.

12522

12520

12542

11708

123/180

11722

12514

12526 12528 12530 12532

OK Cancel Save Delete

Unlink Contract Expand Delete All

12534 12536 12538 12540

12508

United States Patent [19]
Hiratsuka et al.

[11] Patent Number: 5,029,013

[45] Date of Patent: Jul. 2, 1991

[54] APPARATUS FOR EDITING VIDEO TAPES

[75] Inventors: Masaru Hiratsuka, Kanagawa, Japan; Hiroshi Nakashima, Kanagawa, Japan; Kiyoshi Nakagawa, Kanagawa, Japan; Shinya Fukushima, Saitama, Japan

[73] Assignee: Sony Corporation, Tokyo, Japan

[21] Appl.No.: 275,377

[22] Filed: Nov. 23, 1988
Related U.S. Application Data

[30] Foreign Application Priority Data
Nov. 26, 1997. Japan.....62-298604

12518 [1/0]

12510 ↑

12512 ↑

12506

FIG.125

124/180




FIG.126

125/180




FIG.127

126/180

12802

Report Generator

12804

Available Reports:

4.Inventor Employment Information 12806

Levels to Drill Down: 12810

Patents Expiring Between: 12812

Patent Term Years Between: 12814

Execute 12808

Quit

5

01-01-1900

3

12-31-2100

10

And:

And:

FIG.128

127/180




FIG. 129




FIG. 130

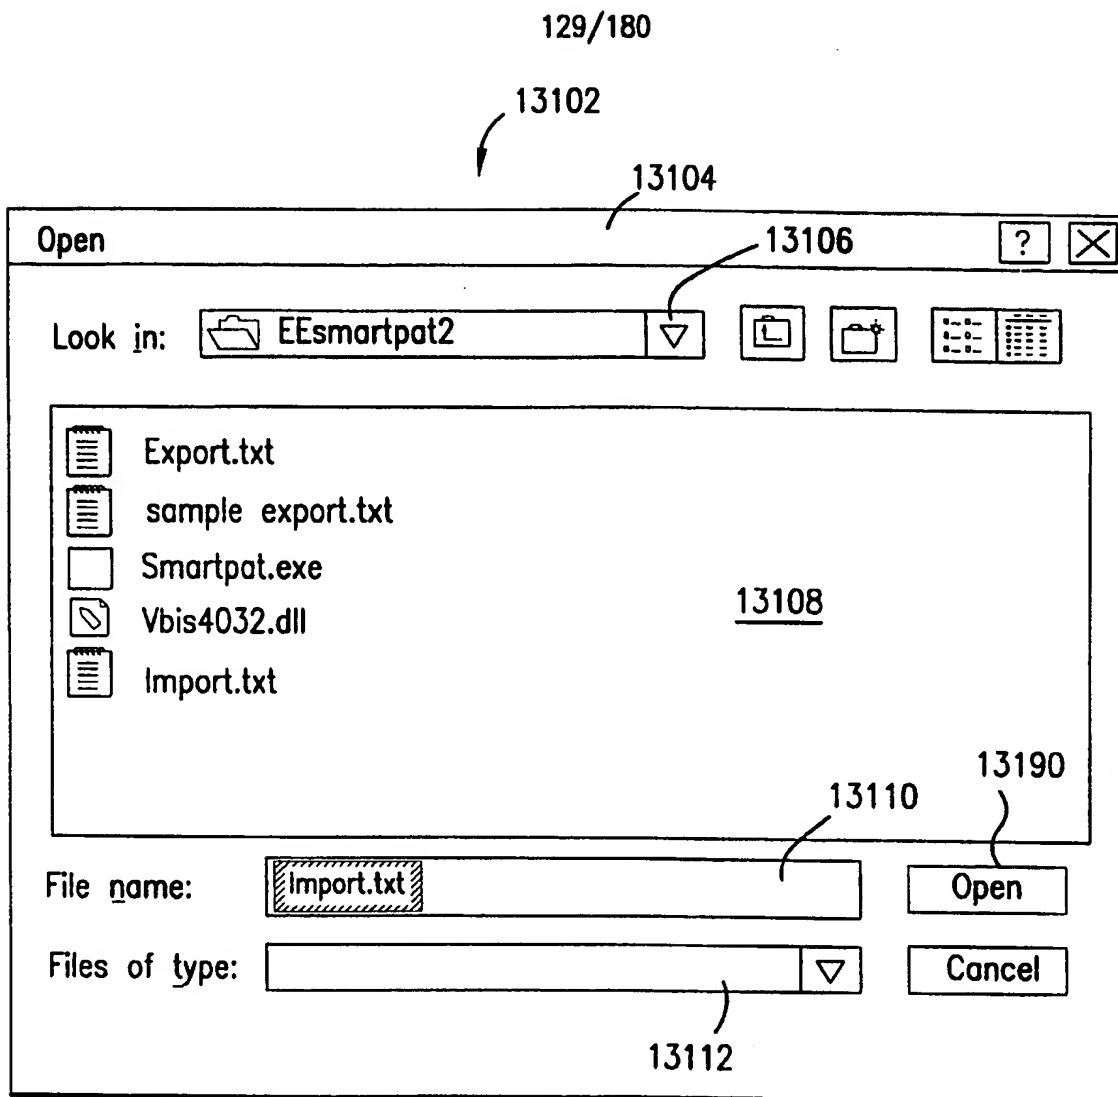


FIG.131





FIG.132

13302



132/180

web.smartpatents.com:8080

13416

13414

13418

13412

11708

13410

13408

11706

11704

13402

13406

FIG. 134

Document#

1

4716476

Title

Topo Recording Apparatus

Assignee

Sony Corpor

Inve

Hiros

Type

C

Purpose of this group

Purpose of this group

Atty. Work Prod.

This group covers all aspects of our analog special effects.

Repository

A Group Mode Remotely

A Potential Patent

Analog Video Editing Systems

Analog Professional Video Editing

Analog Special Effects Contro

Video Color Correction Subsys

Video Reproduction and Copy

Digital Cameras

Disposable Cameras

Internet Appliances

OK

Cancel

Save

Delete




FIG.135




FIG.136




FIG.137

134/180

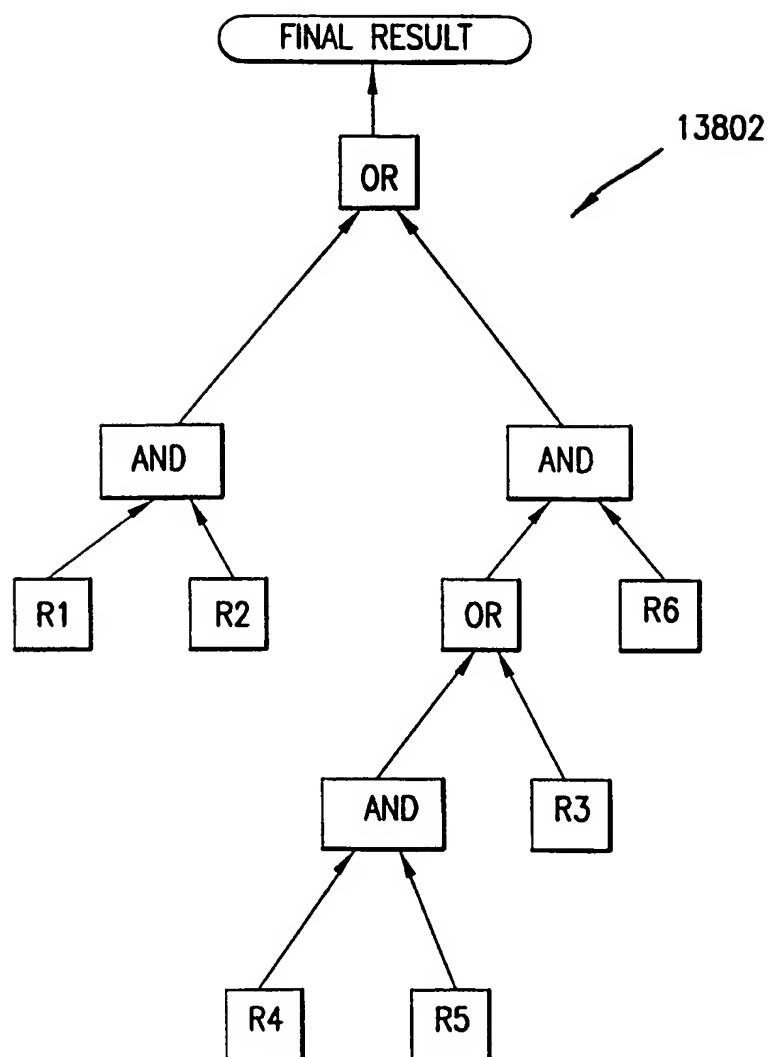


FIG.138

135/180

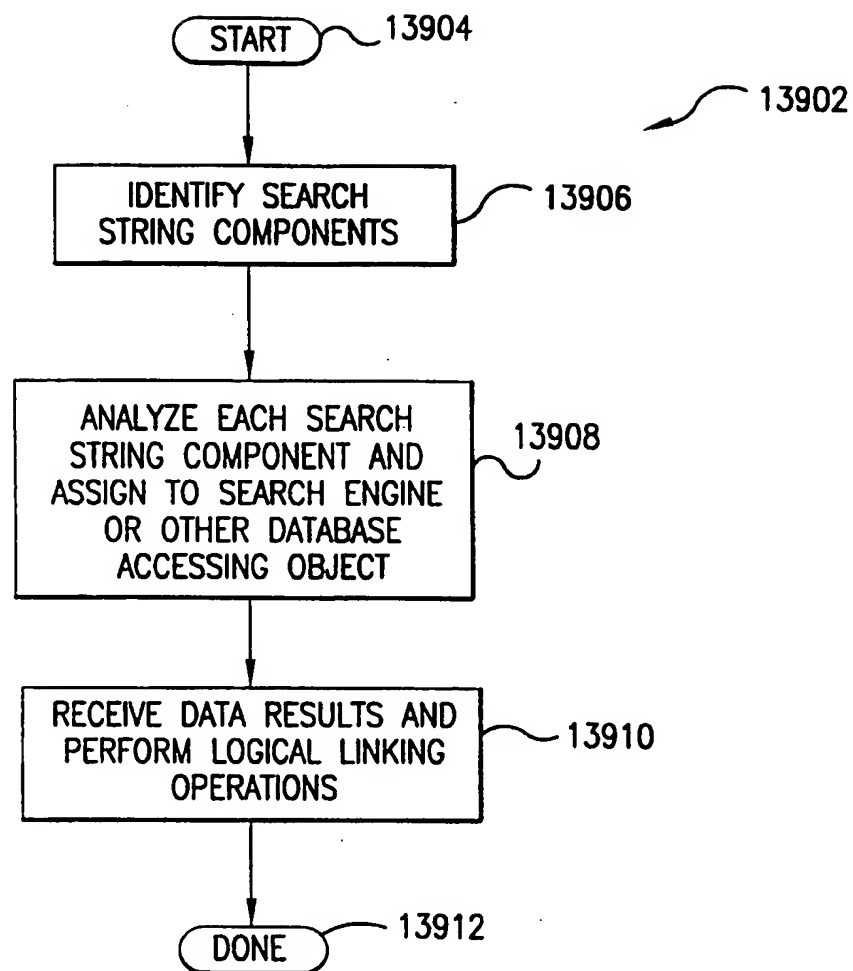



FIG.139

136/180

SmartPatents Patent Search		PAGE 1 OF 1
<div style="display: flex; justify-content: space-between; align-items: center;"><div style="text-align: center;"> SMARTPATENTS</div><div style="text-align: center; flex-grow: 1;"><h2 style="margin: 0;">Patent Search</h2></div></div> <div style="margin-top: 10px;"><p>Search</p><div style="display: flex; justify-content: space-between;"><div style="width: 30%;"><p>Scope of search:</p><p>Patent Number:</p><p>Title:</p><p>Inventor:</p><p>Assignee:</p><p>Class:</p><p>User-Defined Keywords:</p><p>Date of Issue:</p><p>Abstract:</p><p>Full Patent Text:</p></div><div style="width: 65%; border: 1px solid black; padding: 5px;"><div style="display: flex; justify-content: space-between; align-items: center;">All U.S. Patents↓</div><div style="margin-top: 5px;"><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div></div><div style="margin-top: 10px;"><div style="display: flex; align-items: center;"><div style="margin-right: 10px;"><input checked="" type="radio"/> Before <div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div><div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div></div><div style="margin-right: 10px;"><input type="radio"/> Between <div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div> and <div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div></div></div><div style="margin-left: 10px; text-align: right;"><p>(mm/dd/yyyy)</p><p>(mm/dd/yyyy)</p></div></div></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div><div style="border: 1px solid black; height: 20px; width: 100%;"></div></div></div>		

Fields to include in list of search results (Patent Number is always included)

☒ Title

☐ Inventor

☐ User-Defined Keywords

☐ Date of Issue

☐ Assignee

}

14004

Order results by

Score

↓

14006

Patents per page in results

10

(MAX 25)

14008

Search

 14010

[Help with this screen](#)

14002

FIG.140

137/180

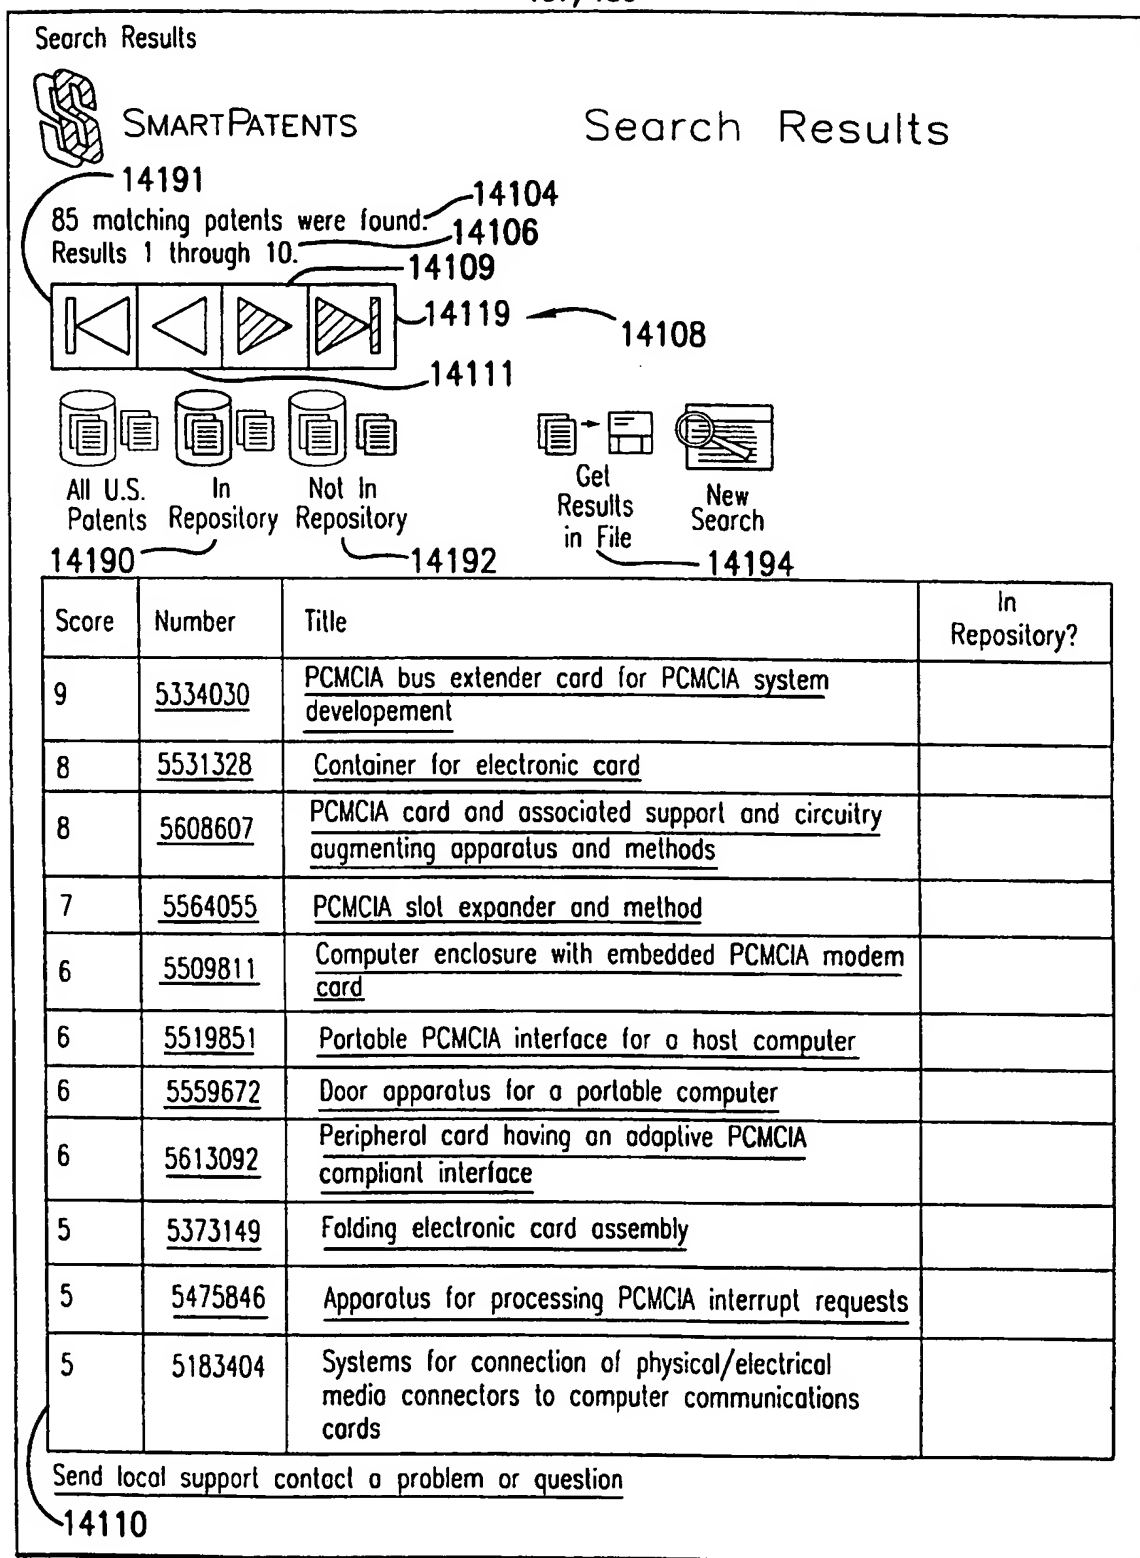


FIG. 141

14102

138/180


Search Results





SMARTPATENTS

Search Results

85 matching patents were found. 14204
Results 11 through 20.



14108

All U.S.
PatentsIn
RepositoryNot In
RepositoryGet
Results
in FileNew
Search

Score	Number	Title	In Repository?
5	<u>5540597</u>	<u>All flex PCMCIA-format cable</u>	X
5	<u>5596562</u>	<u>Controlling method and apparatus for supporting hard disk and/or CD-ROM drives through the PCMCIA interface</u>	
4	<u>5532898</u>	<u>Line current protection circuit for use with a PCMCIA-architecture modem card</u>	
4	<u>5537597</u>	<u>Method and apparatus for supporting real mode card services clients with a protected mode card services implementation</u>	
4	<u>5537654</u>	<u>System for PCMCIA peripheral to execute instructions from shared memory where the system reset signal causes switching between modes of operation by alerting the starting address</u>	
4	<u>5566121</u>	<u>Method for PCMCIA card function using DRAM technology</u>	X
3	<u>5375037</u>	<u>Memory card having a recessed portion with contacts connected to an access card</u>	
3	<u>5486687</u>	<u>Memory card having a recessed portion with contacts connected to an access card</u>	
3	<u>5500517</u>	<u>Apparatus and method for data transfer between stand alone integrated circuit smart card terminal and remote computer of system operator</u>	
3	<u>5537558</u>	<u>Apparatus and method for communicating multiple devices through one PCMCIA interface</u>	

Send local support contact a problem or question

FIG. 142

14202

139/180

Search Results




SMARTPATENTS

Search Results

85 matching patents were found. 14304
Results 21 through 30.



14108

All U.S.
PatentsIn
RepositoryNot In
RepositoryGet
Results
in FileNew
Search

14310

Score	Number	Title	In Repository?
3	<u>5555510</u>	<u>Automatic computer card insertion and removal algorithm</u>	
3	<u>5611055</u>	<u>Method and apparatus for implementing a PCMCIA auxiliary port connector for selectively communicating with peripheral devices</u>	
3	<u>5615328</u>	<u>PCMCIA SRAM card function using DRAM technology</u>	X
2	<u>5436621</u>	<u>Messaging peripheral with secure message data function</u>	
2	<u>5455572</u>	<u>Selective call receiver with computer interface message notification</u>	
2	<u>5475563</u>	<u>PCMCIA card heat removal apparatus and methods</u>	
2	<u>5481616</u>	<u>Plug-in sound accessory for portable computers</u>	
2	<u>5485505</u>	<u>Apparatus and method for remotely initiating operation of a cellular telephone</u>	
2	<u>5488572</u>	<u>Portable computer system for docking to an expansion base unit</u>	
2	<u>5526226</u>	<u>Information processing apparatus and device accomodated therein with a mechanism for locking a cover</u>	X

Send local support contact a problem or question

14302

FIG.143

140/180

Patent Abstract: 5183404



SMARTPATENTS

Patent Abstract

New
Search

14404

Patent Number	5183404
Title	Systems for connection of physical/electrical media connectors to computer communications cards
Issue Date	02/02/93
Class(es)	439/676 439/372 439/329 439/131 439/55
Inventor(s)	Guy M. Dake, Sandy, UT Stephen C. Aldous, Salt Lake City, UT
Assignee	Megahertz Corporation, Salt Lake City, UT

Abstract:

14406

A communications card capable of being mounted in electrical communications with a computer has formed therethrough an aperture so sized and shaped as to be capable of receiving a physical/electrical media connector. The media connector has a biased retention clip, a contact pin block, and contact pins. The retention clip has several standardized characteristics including a broad fixed end protruding from an outer surface of the contact pin block. The broad fixed end tapers abruptly at a transition notch down to a narrow free end, capable of being manipulated by a user to remove the physical/electrical media connector from the aperture in the communications card. In use, a media connector is inserted directly into the aperture in the communications card, the aperture being in contact with a plurality of contact wires fixed within the communications card. The communications card is divided into a retractable access portion of the communications card which can be directly accessed by manipulating an actuating mechanism releasing a retention means thereby allowing a spring to push the retractable access portion of the card outside of the computer housing. The retractable access portion of the communications card may be reinserted back into the computer housing to be carried internally when not in use.


FIG. 144

14402

141/180

Patent Text: 5615328

Page 1 of 15



BIBLIOGRAPHY — 14512
PARENT CASE TEXT — 14514
CROSS REFERENCE TO RELATED APPLICATIONS — 14516
BRIEF DESCRIPTION OF THE DRAWINGS — 14518
DETAILED DESCRIPTION OF THE INVENTION — 14520
CLAIMS — 14522

} 14524

United States Patent [19]

Hodderman et al.

[11] Patent Number: 5,615, 328

[45] Date of Patent: Mar. 25, 1997

[54] PCMCIA SRAM CARD FUNCTION USING DRAM TECHNOLOGY

[75] Inventors: Scott J. Hodderman, Pleasant Valley, NY;
Kraig R. White, Colchester, VT[73] Assignee: International Business Machines Corporation
, Armonk, NY

[21] Appl. No.: 521,509

[22] Filed: Aug. 30, 1995

Related U.S. Application Data

[51] Int. Cl.6 G01R 31/28; G06F 11/00

[52] U.S. Cl. 395/182.2; 365/222

[58] Field of Search 395/182.2; 365/222

[56] References Cited

U.S. PATENT DOCUMENTS

4,316,248 2/1982 Retter 395/182.2

14526

14502

FIG. 145A

142/180

Patent Text: 5615328

Page 2 of 15

	4,451,742	5/1984 Aswell	395/182.2 X	
	4,710,903	12/1987 Hereth et al.	365/194	
	4,766,567	8/1988 Kato	395/182.2 X	
14536	4,958,322	9/1990 Kosugi et al. ..	365/189.01	
	5,206,830	4/1993 Isobe	365/194	
	5,208,779	5/1993 Walther	365/222	14528
	5,229,970	7/1993 Lee	365/222	
	5,276,843	1/1994 Tillinghast et al. ...	395/425	
	5,283,905	2/1994 Soodeh	395/182.2 X	
	5,289,424	2/1994 Ito et al.	365/222	
	5,333,128	7/1994 Yoon	365/222	
	5,335,201	8/1994 Walther	365/222	
	5,414,861	5/1995 Horning	395/182.2 X	
	5,430,681	7/1995 Sugawara	365/222	14538

OTHER PUBLICATIONS

Brown, New IC Stacking Process Ideal For High-Density Memory Module and Hybrid Applications, Electro/95 Int'l Professional Program Proceedings, Apr. 1995, pp. 3-16.

Primary Examiner--Robert W. Beausoliel, Jr.
 Assistant Examiner--Glenn Snyder
 Attorney, Agent, or Firm--Marc A. Eherlich

[57]

ABSTRACT

An apparatus may be used with a computer system having a PCMCIA interface. The apparatus employs a DRAM device and logic for converting the PCMCIA SRAM control signals into DRAM control signals, so as to permit the communication of data and control signals between the computer system and the DRAM device. The apparatus further provides controls for refreshing the DRAM device, and for arbitrating between the functions of refreshing the DRAM and providing for communication between the DRAM and the computer system. The apparatus further provides the power management functions required for operating a DRAM device in a PCMCIA environment.

11 Claims, 6 Drawing Sheets

PCMCIA SRAM CARD FUNCTION USING DRAM
 TECHNOLOGY

CROSS REFERENCE TO RELATED APPLICATIONS

The present application related to the following co-pending patent applications:

PCMCIA voltage loss detection, U.S. patent application
 Ser. No. 08/521,498, by Scott J. Hadderman

A Method for PCMCIA SRAM Card Function Using

FIG. 145B

143/180

Patent Text: 5615328

Page 3 of 15

DRAM Technology, U.S. patent application Ser. No. 08/ 521,508, by Scott J. Hadderman and
Kraig R. White

14534

These co-pending applications and the present application are owned by one and the same assignee, International Business Machines Corporation of Armonk, N.Y.

The descriptions set forth in these co-pending applications are hereby incorporated into the present application by this reference.

FIELD OF THE INVENTION

This invention is generally related to computer systems and particularly to computer systems utilizing a PCMCIA SRAM interface for PCMCIA memory cards.

GLOSSARY OF TERMS

While dictionary meanings are also implied by certain terms used here, the following glossary of some terms may be useful. DRAM Dynamic Random Access Memory.

SRAM Static Random Access Memory.

PCMCIA Personal Computer Memory Card International Association.

ASIC Application Specific Integrated Circuit.

Host Computer System Computer to which the PCMCIA memory card is currently attached.

REFERENCES USED IN THE DISCUSSION OF THE INVENTION

During the detailed description which follows the following work will be referenced as an aid for the reader. The additional reference is:

PCMCIA PC CARD STANDARD Release 2.1 Personal Computer Memory Card International Association 1030 East Duane Avenue Suite G Sunnyvale, Calif. 94086

This additional reference is incorporated by reference.

BACKGROUND OF THE INVENTION

The PCMCIA specification was created to standardize a removable peripheral device for personal computers. The PCMCIA specification describes the features required for PCMCIA devices. These features include physical dimensions and electrical specifications. The PCMCIA electrical specification can further be divided into two card categories, input/output cards and memory cards. This invention is described by the memory card category of the PCMCIA specification.

Although the memory specification is intended to describe PCMCIA devices using general memory technology, the specification tends to limit memory selection to either SRAM or FLASH technology. The specification

FIG.145C

144/180

Patent Abstract: 5206830



SMARTPATENTS

Patent Abstract

New
Search

Patent Number	5206830
Title	Refresh control circuit of pseudo static random access memory and pseudo static random access memory apparatus
Issue Date	04/27/93
Class(es)	307/606 307/602 307/596 307/592 365/233 365/222 365/194
Inventor(s)	Hisashi Ueno, Kawasaki, Japan Mitsuo Isobe, Tama, Japan
Assignee	Kabushiki Kaisha Toshiba, Kawasaki, Japan

Abstract:

A refresh control circuit for a pseudo static random access memory includes a refresh control signal output circuit for outputting a refresh control signal to accomplish refresh control of the pseudo static random access memory, and includes a delay circuit. A first chip enable signal from a control device such as a MPU is delayed by the delay circuit and outputted as a second chip enable signal for the PSRAM. As the first chip enable signal level changes from a selection level to a non-selection level, the refresh control signal level changes to a non-refresh level. This state is maintained for a predetermined period. After the second chip enable signal changes from the selection level to the non-selection level, the refresh control signal returns from the non-refresh level to the refresh level. Thus, the PSRAM enters into the refresh state during the non-selection state, and is refreshed. This refresh operation is necessarily performed after an access to PSRAM.

Order this patent from SmartPotents


14604

FIG. 146

14602


145/180

Search Results Page 1 of 1




Search Results: Patents in Repository


7 matching patents were found.
Results 1 through 7.




All U.S.
Patents




In
Repository




Not In
Repository



Get
Results
in File



Skim
Images



New
Search

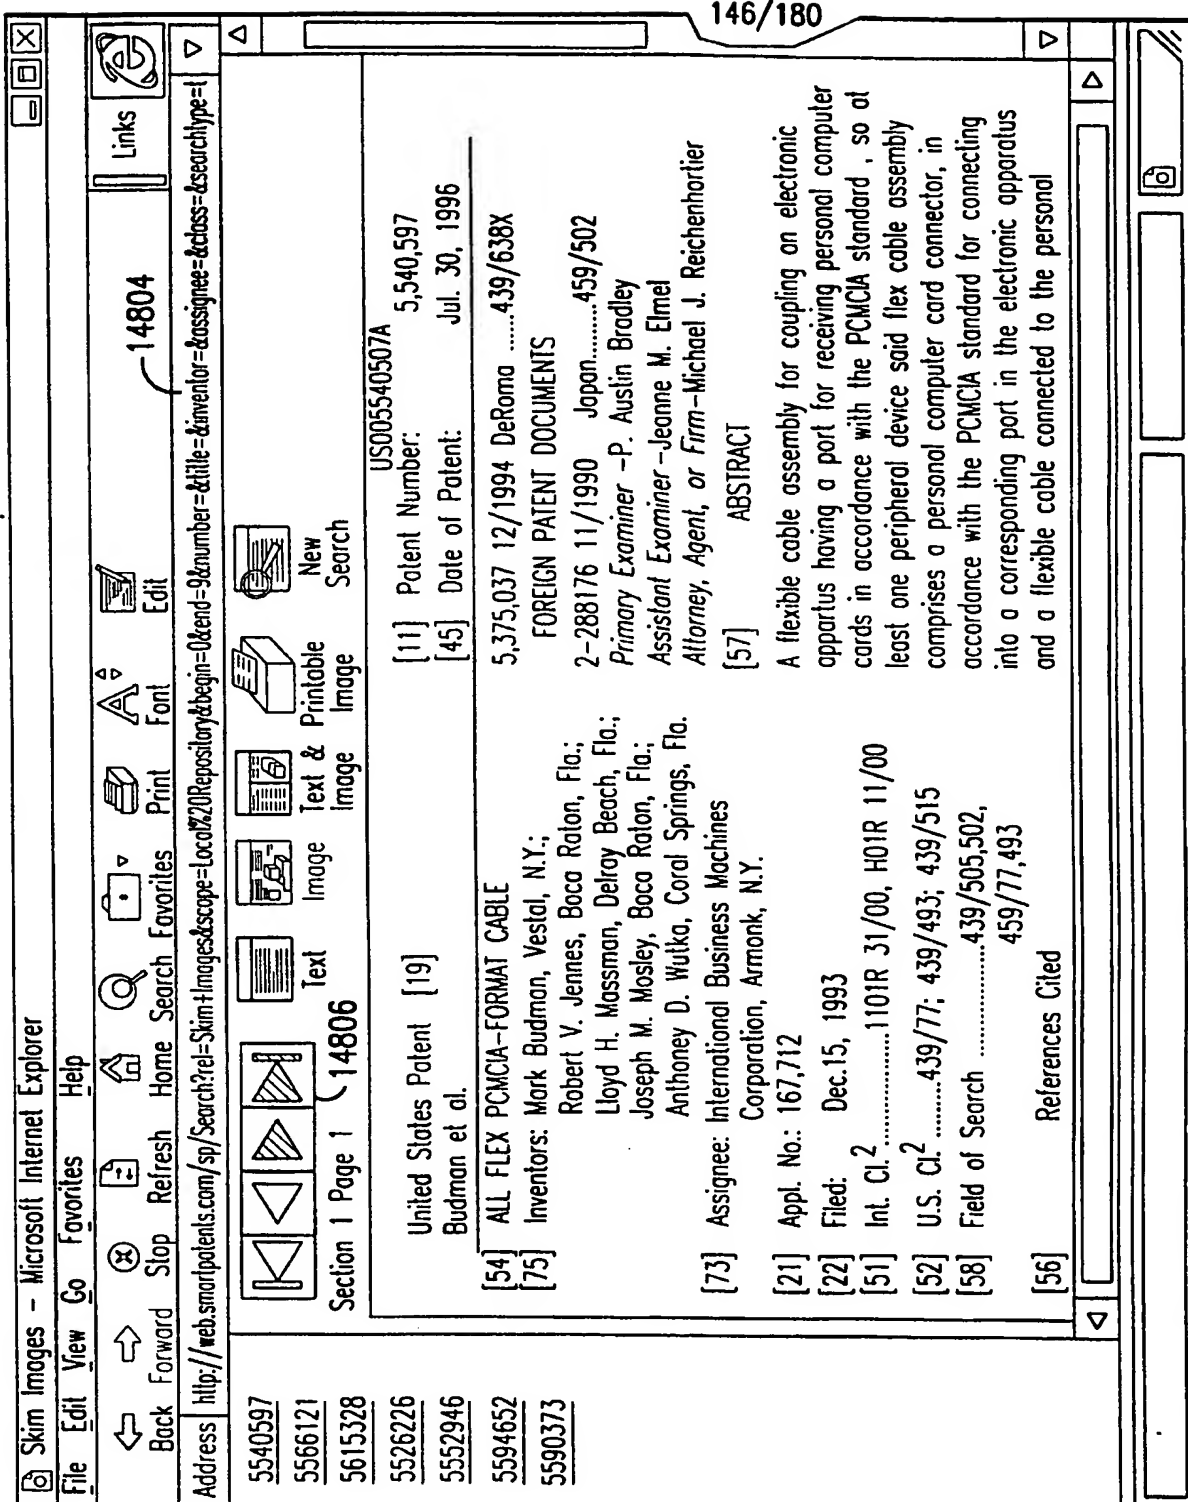
14704 14706

Score	Number	Title
5	<u>5540597</u>	<u>All flex PCMCIA-format cable</u>
4	<u>5566121</u>	<u>Method for PCMCIA card function using DRAM technology</u>
3	<u>5615328</u>	<u>PCMCIA SRAM card function using DRAM technology</u>
2	<u>5526226</u>	<u>Information processing apparatus and device accommodated therein with a mechanism for locking a cover</u>
2	<u>5552946</u>	<u>Compliant rail for shock protection of a PCMCIA DASD</u>
2	<u>5594952</u>	<u>Device for integration between voice and data radio communication</u>
1	<u>5590373</u>	<u>Field programming apparatus and method for updating programs in a personal communications device</u>

[Send local support contact a problem or question](#)

14702

FIG.147



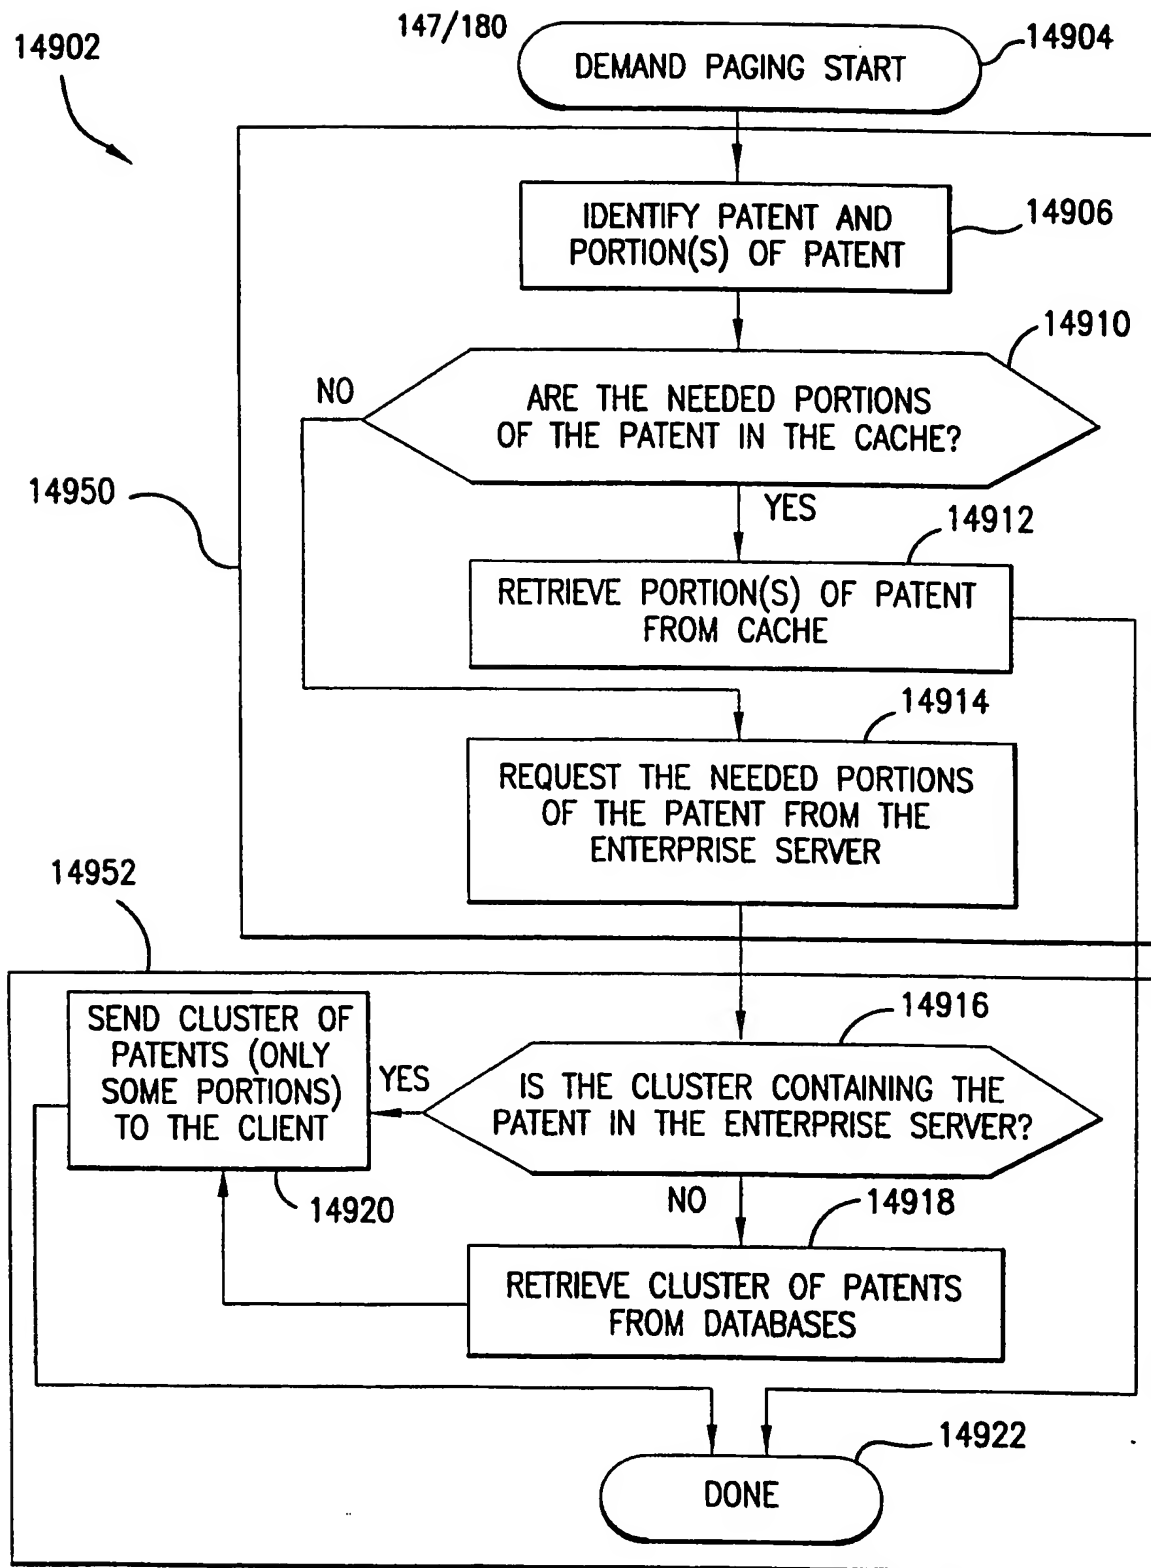


FIG.149




FIG.150




FIG.151




FIG.153

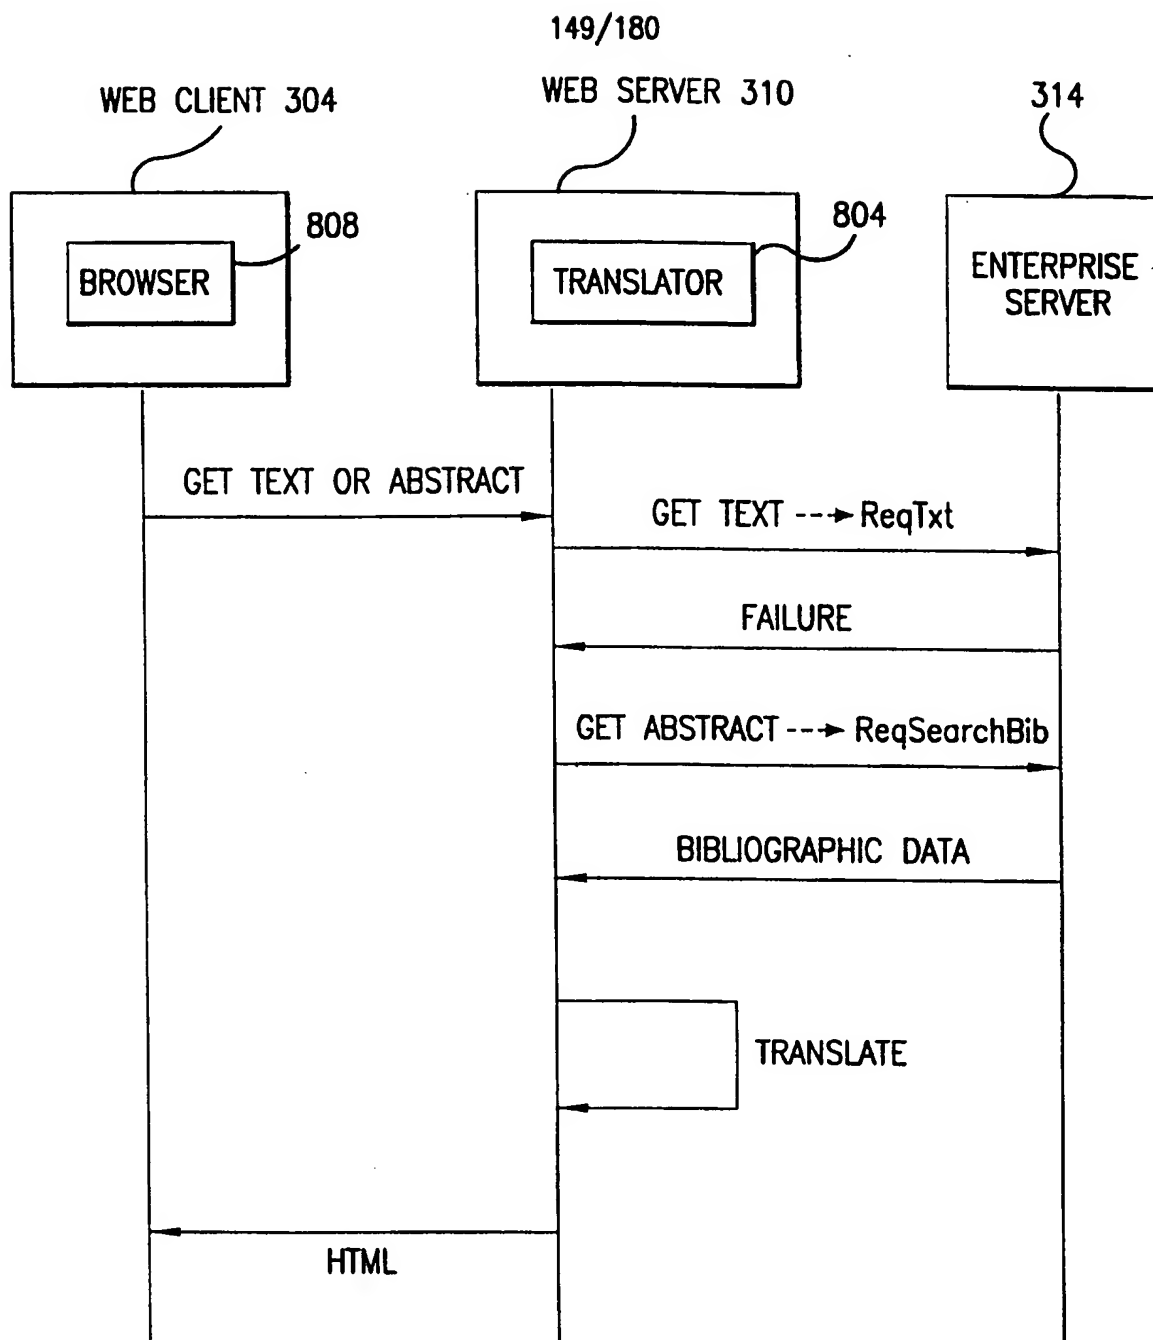


FIG.152

SmartPatent(R) Workbench Enterprise Edition-(Myron 8000)																																									
File		Edit	Document	Image	Annotation	Window	Help																																		
<div style="display: flex; justify-content: space-between;"> <div> </div> <div> </div> </div>																																									
<div style="display: flex;"> <div style="flex: 1;"> <p> Repository</p> <ul style="list-style-type: none"> + A New Group + A New Group - a new name - Bravo + Charlie Project - Chuck's group test + Dave's "Test" Group - Division A + F - For Dan's Test + Foxtrot Project, fox, fox + Garbage + Gordon's Test Hierarchy + hello - Hotel Project - India - Juliet Project - Kilo Project - Lima Project - Mike Project </div> <div style="flex: 1;"> <div style="border: 1px solid black; height: 20px; margin-bottom: 5px;"></div> <table border="1"> <thead> <tr> <th></th> <th>Document#</th> <th>Title</th> <th>Assignee</th> <th>Inventor</th> <th>Class</th> <th>Issue Date</th> <th>Keywords</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>4088268</td> <td></td> <td></td> <td></td> <td></td> <td></td> <td>Y</td> </tr> <tr> <td>2</td> <td>4103314</td> <td></td> <td></td> <td></td> <td></td> <td></td> <td>Y</td> </tr> <tr> <td>3</td> <td>4733336</td> <td>LightedInfor</td> <td>Donnelly Cor</td> <td>Brian I. Scog</td> <td>362/142--</td> <td>03/22/86</td> <td>Y</td> </tr> </tbody> </table> </div> </div>											Document#	Title	Assignee	Inventor	Class	Issue Date	Keywords	1	4088268						Y	2	4103314						Y	3	4733336	LightedInfor	Donnelly Cor	Brian I. Scog	362/142--	03/22/86	Y
	Document#	Title	Assignee	Inventor	Class	Issue Date	Keywords																																		
1	4088268						Y																																		
2	4103314						Y																																		
3	4733336	LightedInfor	Donnelly Cor	Brian I. Scog	362/142--	03/22/86	Y																																		
<div style="display: flex; justify-content: space-between;"> <div> </div> <div> </div> </div>																																									
<div style="display: flex;"> <div style="flex: 1;"> <p> Repository</p> <ul style="list-style-type: none"> + A New Group + A New Group - a new name - Bravo + Charlie Project - Chuck's group test + Dave's "Test" Group - Division A + F - For Dan's Test + Foxtrot Project, fox, fox + Garbage + Gordon's Test Hierarchy + hello - Hotel Project - India - Juliet Project - Kilo Project - Lima Project - Mike Project </div> <div style="flex: 1;"> <div style="border: 1px solid black; height: 20px; margin-bottom: 5px;"></div> <table border="1"> <thead> <tr> <th></th> <th>Document#</th> <th>Title</th> <th>Assignee</th> <th>Inventor</th> <th>Class</th> <th>Issue Date</th> <th>Keywords</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>4088268</td> <td></td> <td></td> <td></td> <td></td> <td></td> <td>Y</td> </tr> <tr> <td>2</td> <td>4103314</td> <td></td> <td></td> <td></td> <td></td> <td></td> <td>Y</td> </tr> <tr> <td>3</td> <td>4733336</td> <td>LightedInfor</td> <td>Donnelly Cor</td> <td>Brian I. Scog</td> <td>362/142--</td> <td>03/22/86</td> <td>Y</td> </tr> </tbody> </table> </div> </div>											Document#	Title	Assignee	Inventor	Class	Issue Date	Keywords	1	4088268						Y	2	4103314						Y	3	4733336	LightedInfor	Donnelly Cor	Brian I. Scog	362/142--	03/22/86	Y
	Document#	Title	Assignee	Inventor	Class	Issue Date	Keywords																																		
1	4088268						Y																																		
2	4103314						Y																																		
3	4733336	LightedInfor	Donnelly Cor	Brian I. Scog	362/142--	03/22/86	Y																																		

(CONT. ON FIG.154B)

FIG. 154A

151/180

(CONT. FROM FIG. 154A)




FIG. 154B

SmartPatent(R) Workbench Enterprise Edition--(Myron 0000)

File Edit Document Image Note Window Help

Document #	Title	Assignee	Inventor	Class	Issue Date	Keywords
1	5240720 Dog chew w		Herbert R. A	426/2---	08/31/93	N
2	0336693 Dog Training		Dan Miller, A	030/199---	06/22/93	N

Type Title

Repository 15506

- + A New Group
- + A New Group 11724
- + a new name
- + Bravo
- + Charlie Project
- + Chuck's group test
- + Dave's "Test" Group
- + Division A
- + F
- + For Dan's Test
- + Foxtrot Project, fox, fox
- + Garbage
- + Gordon's Test Hierarchy
- + hello
- + Hotel Project
- + India
- + Juliet Project
- + Kilo Project
- + Lima Project
- + Mike Project

(CONT. ON FIG.155B)

FIG. 155A

(CONT. FROM FIG.155A)




FIG.155B

154/180

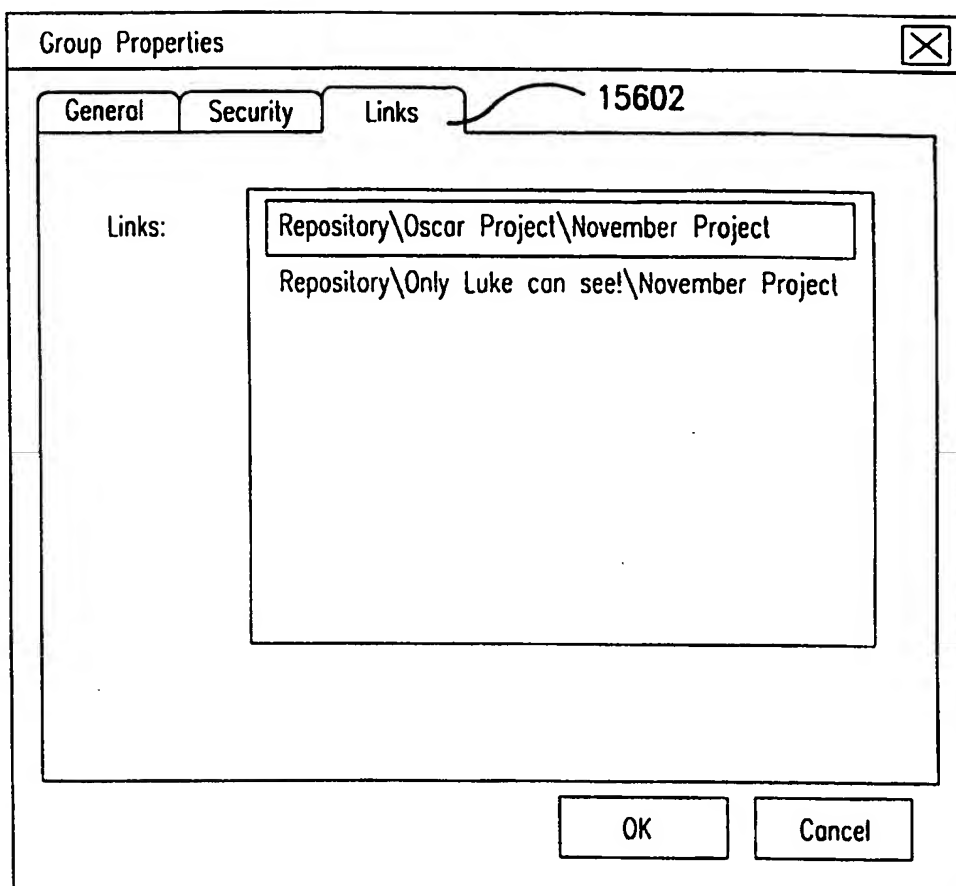


FIG.156

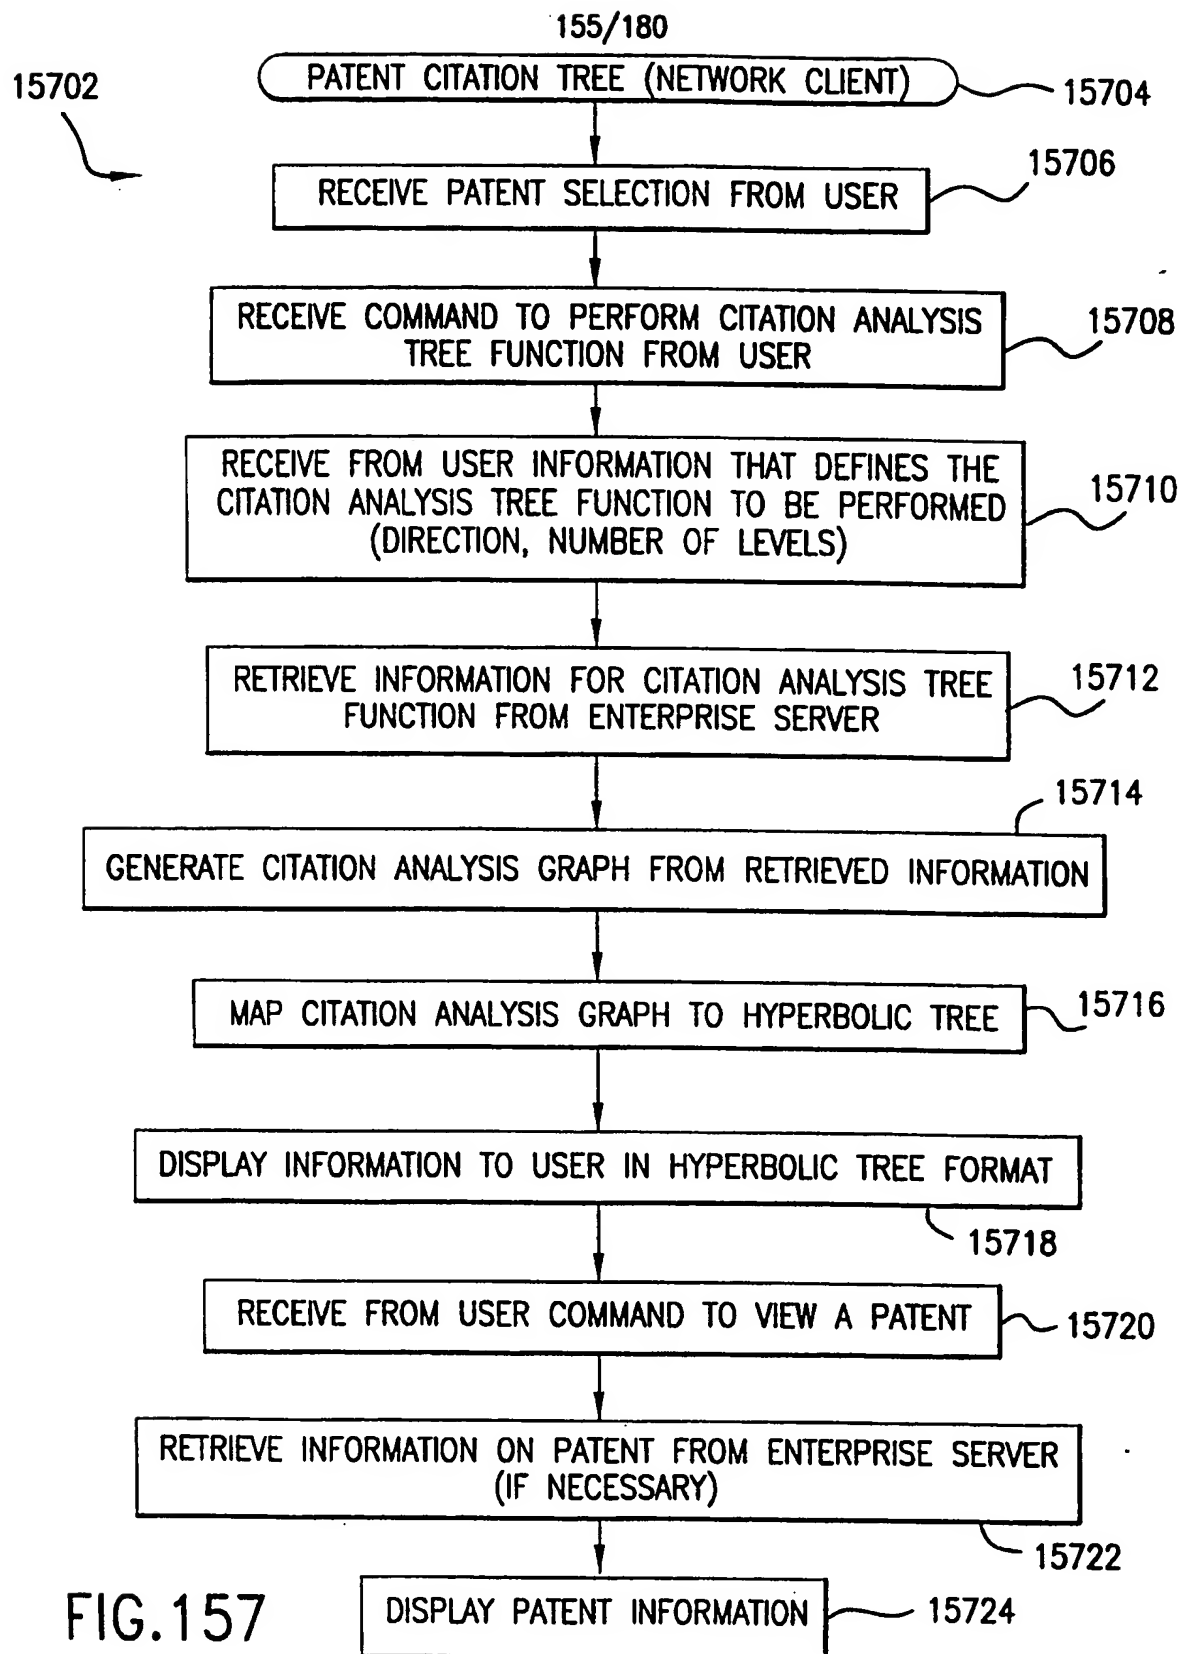


FIG.157

156/180

15712

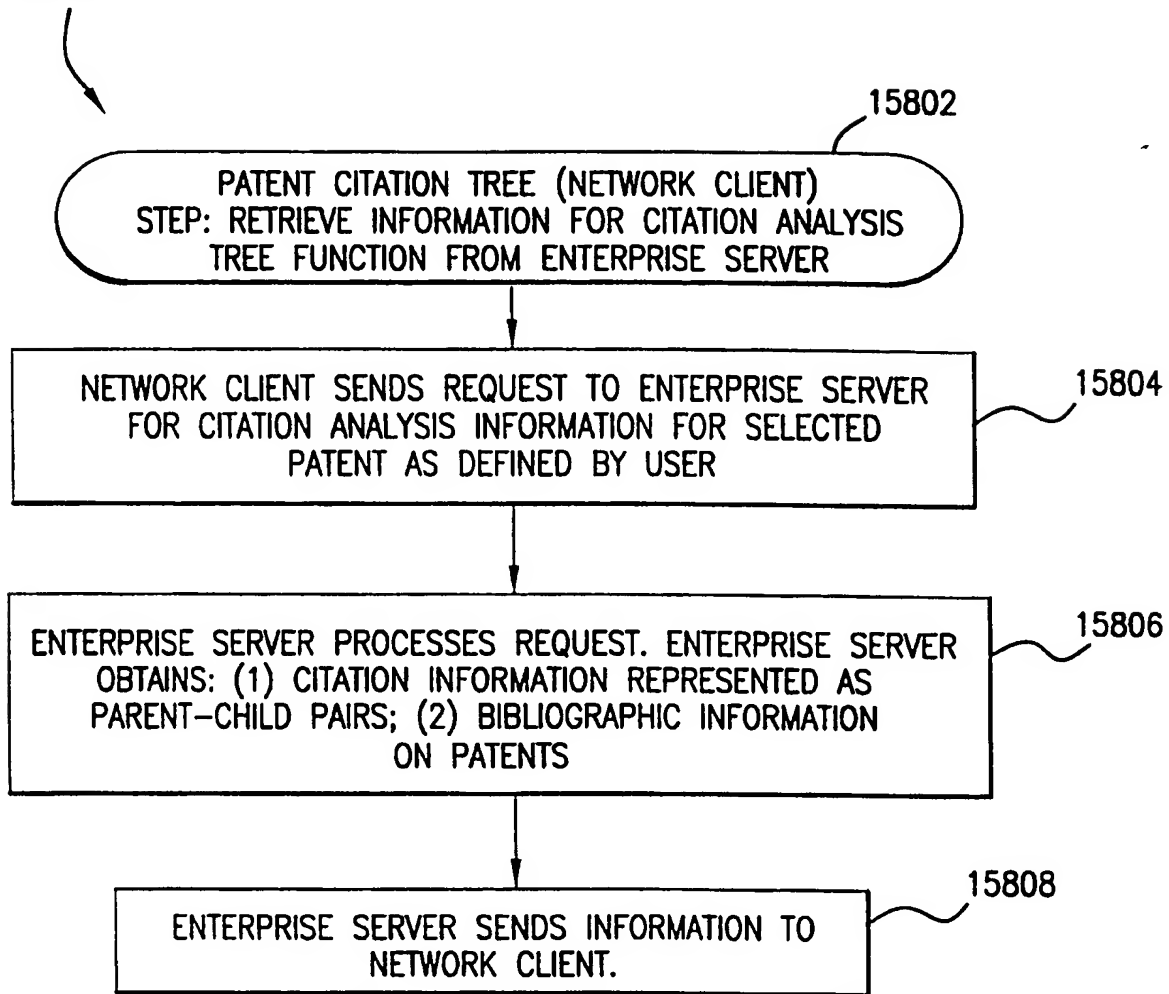


FIG.158

157/180

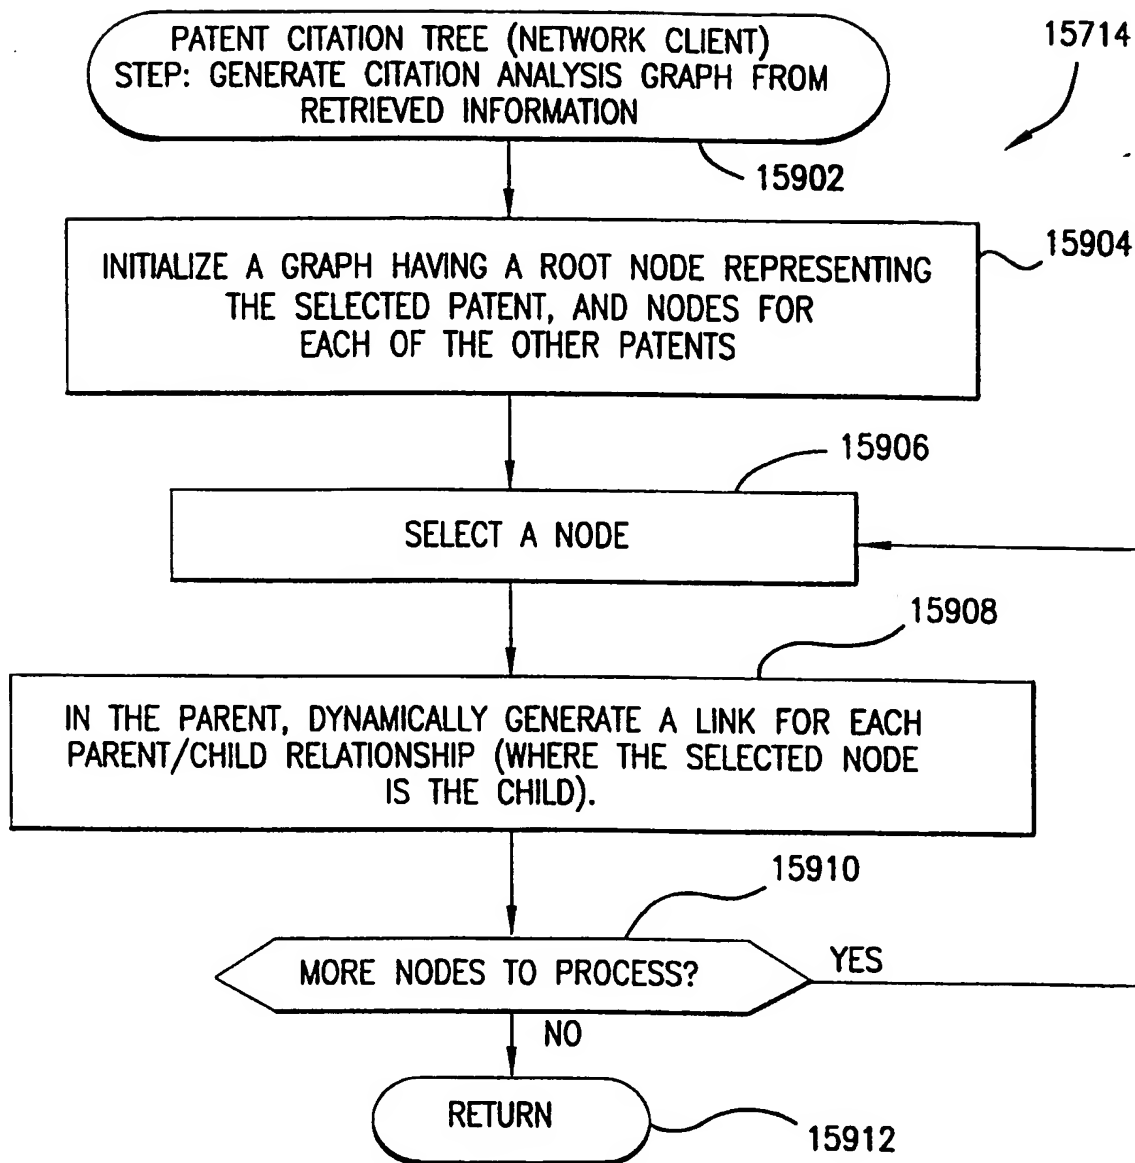


FIG.159

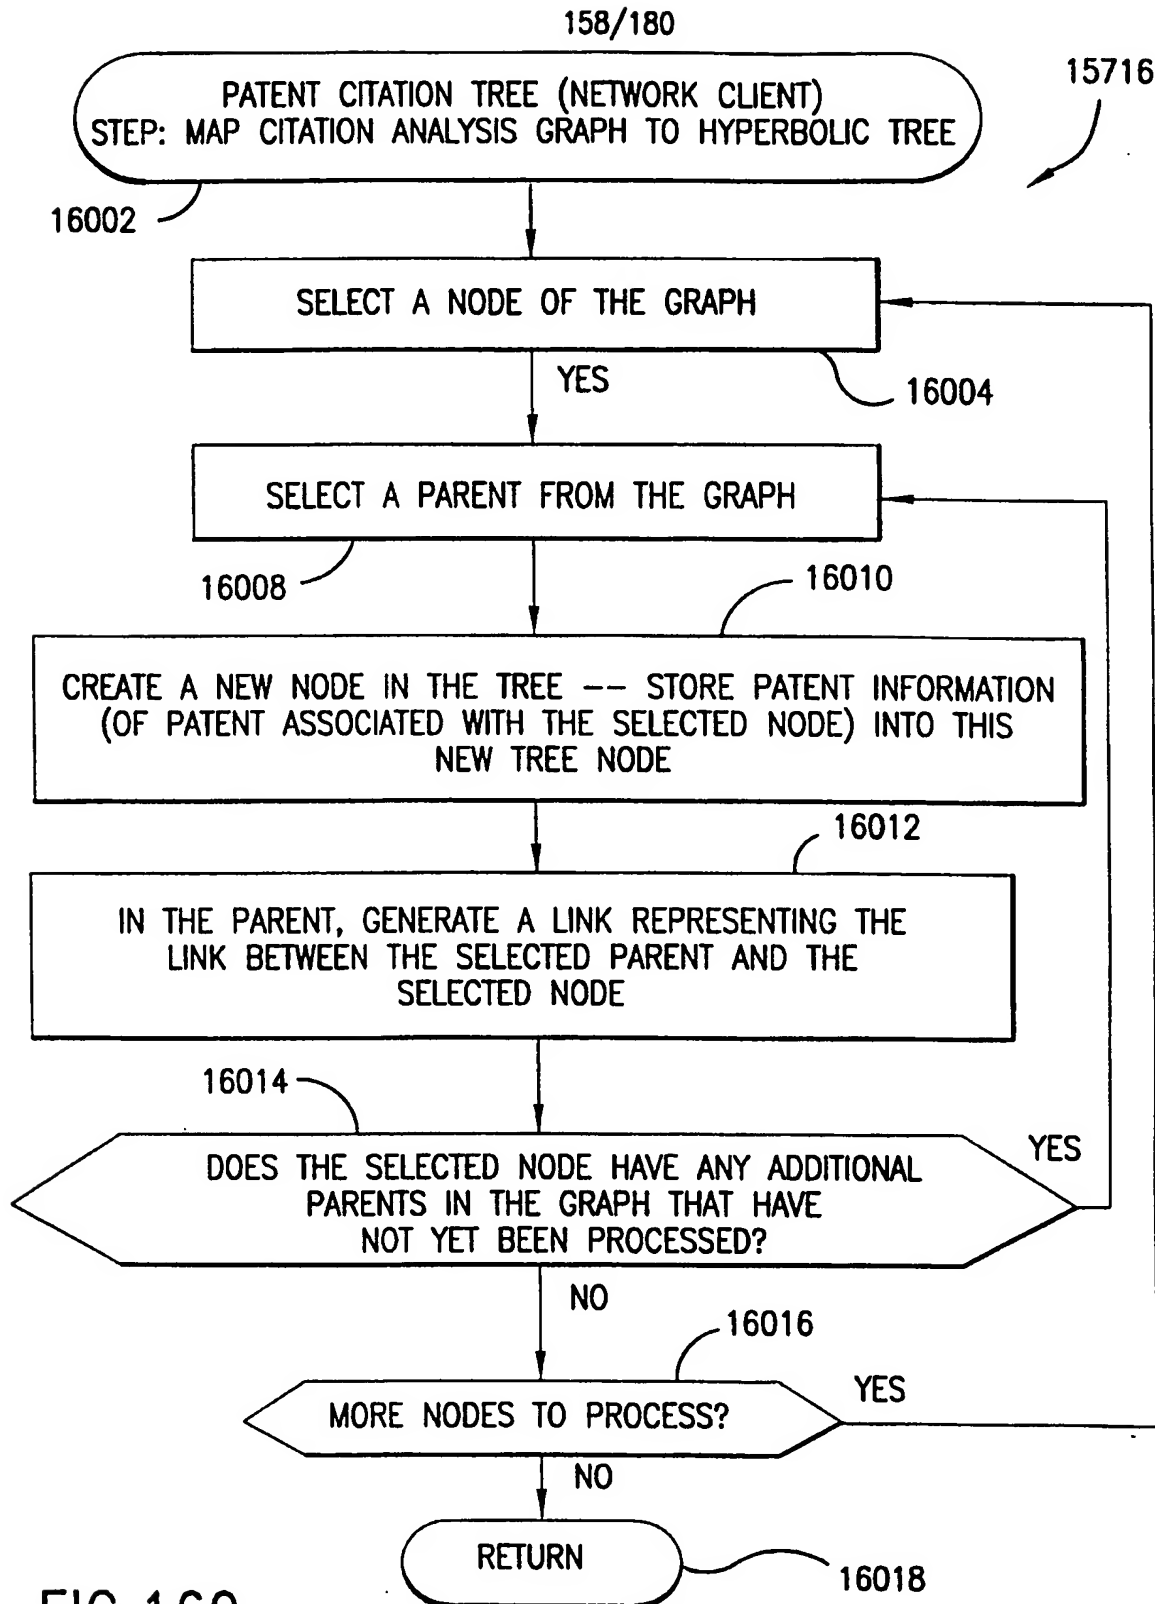


FIG.160

159/180

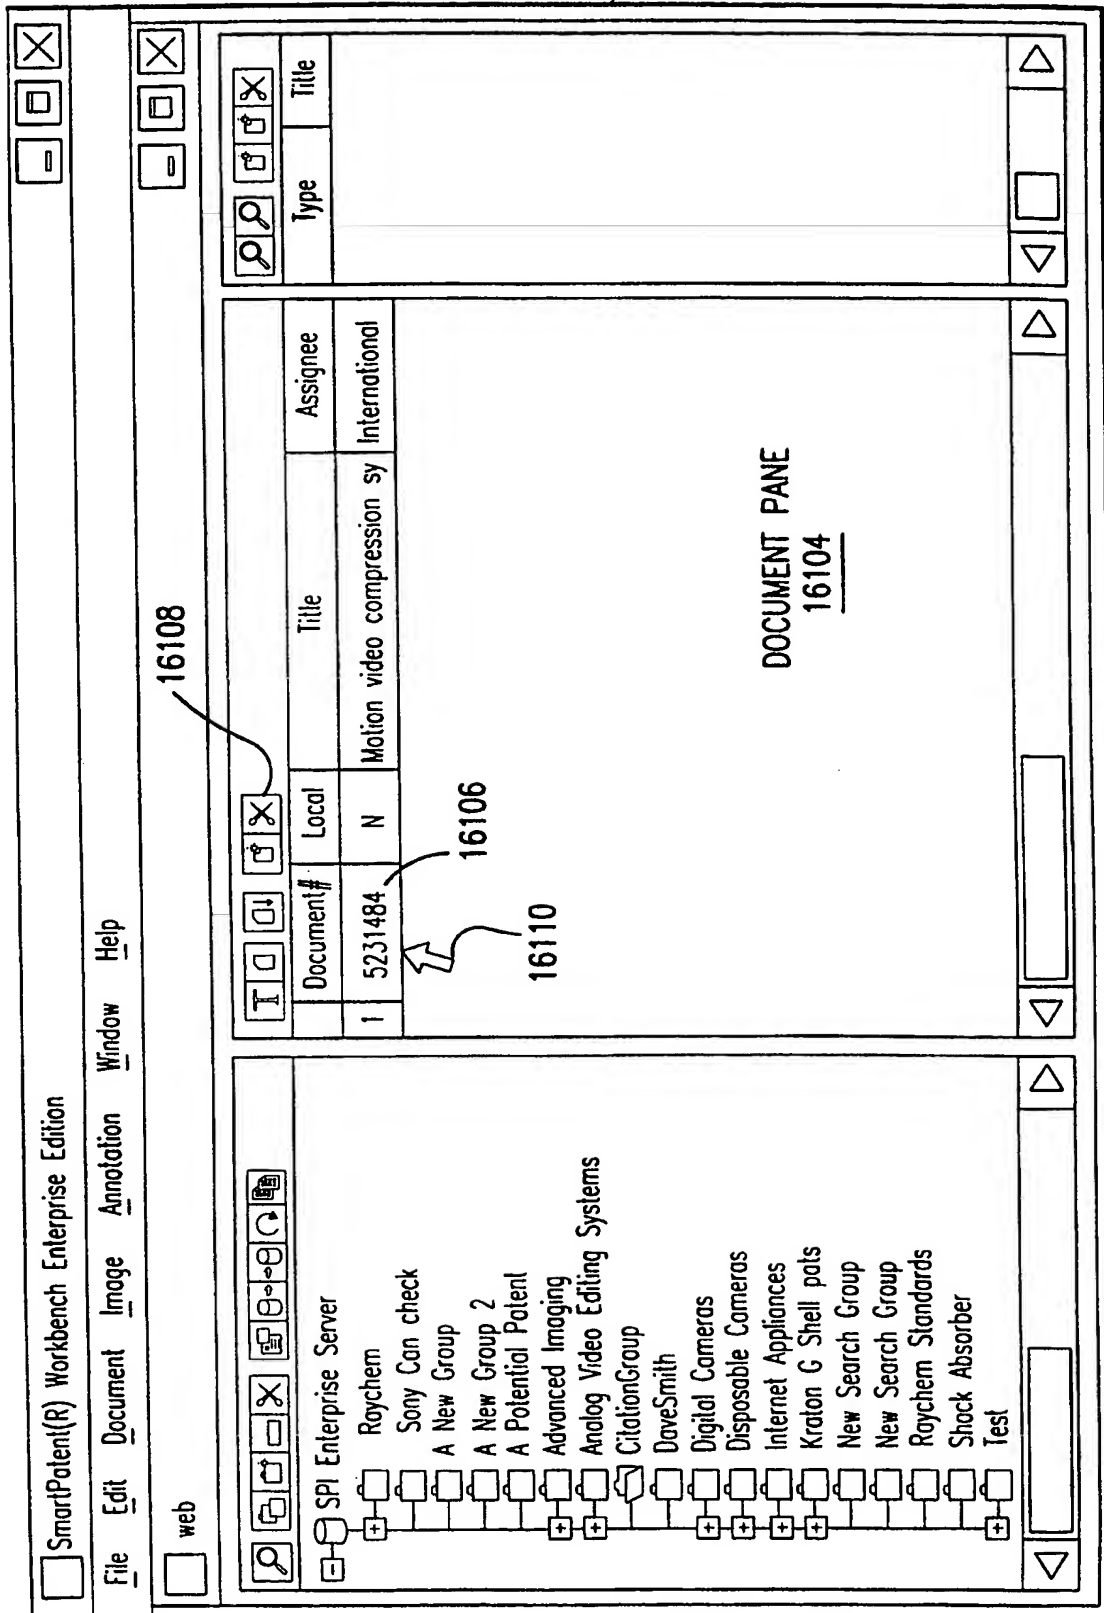


FIG. 161

160/180




FIG. 162

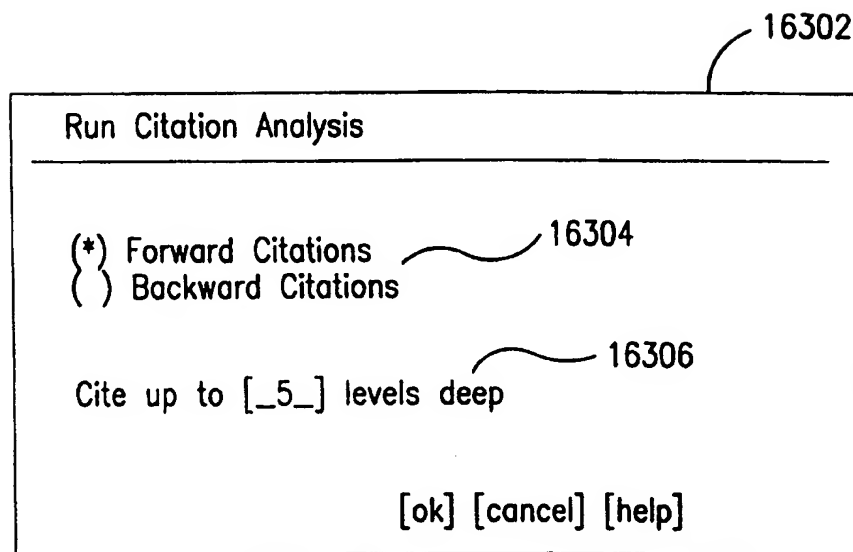



FIG. 163

161/180



CONSOLE
16402

FIG.164

162/180

16502

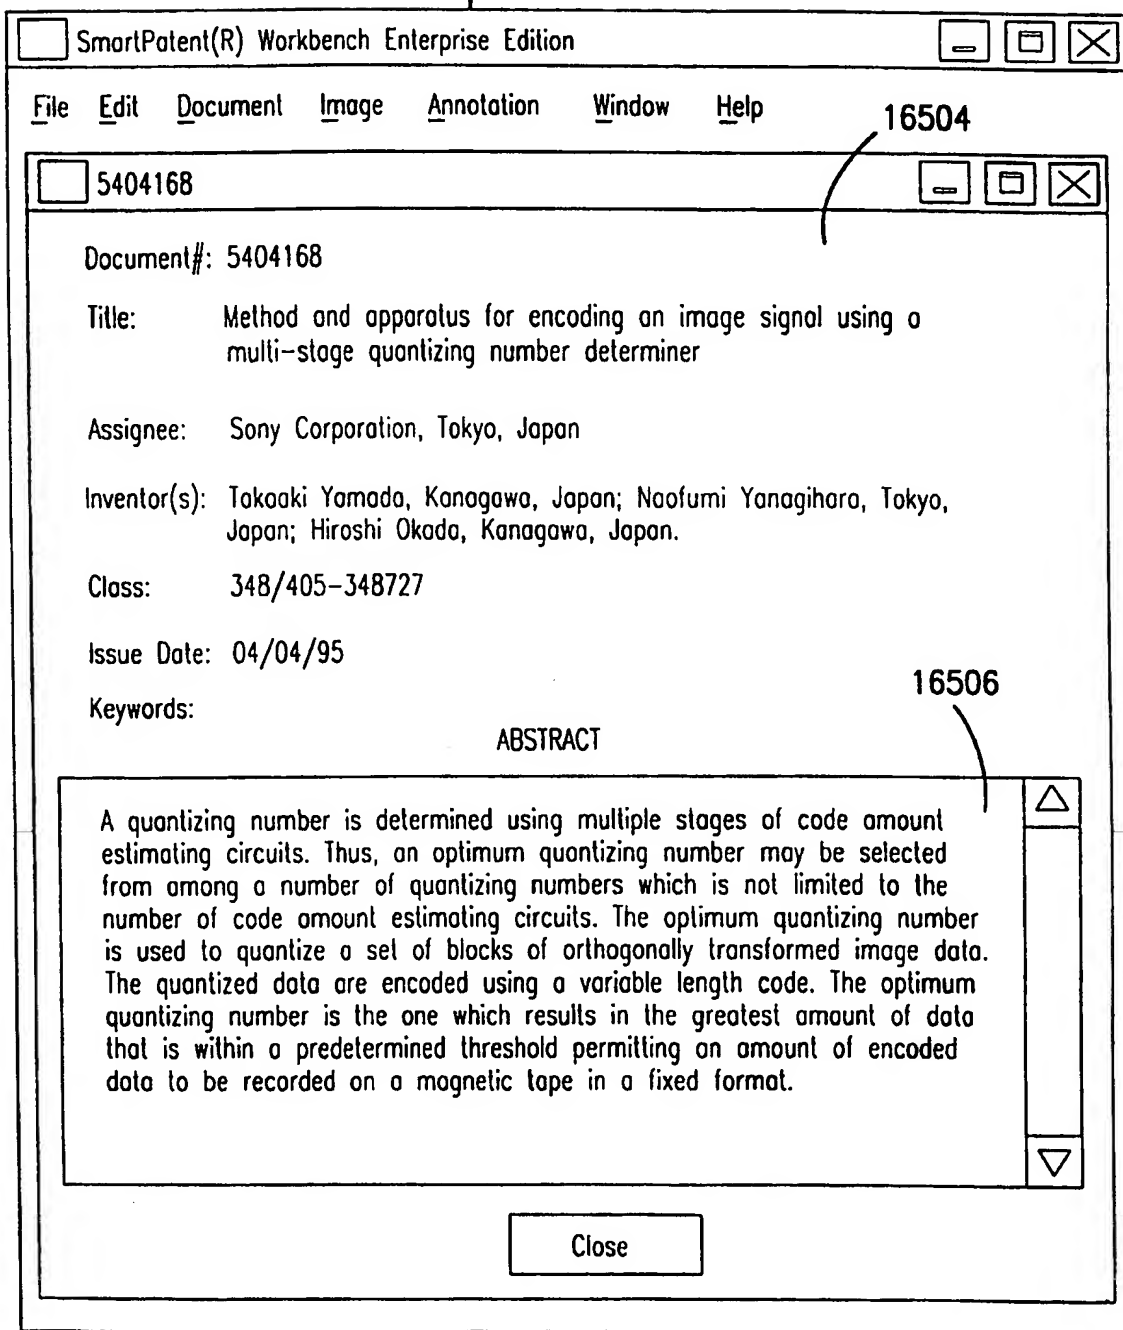


FIG. 165

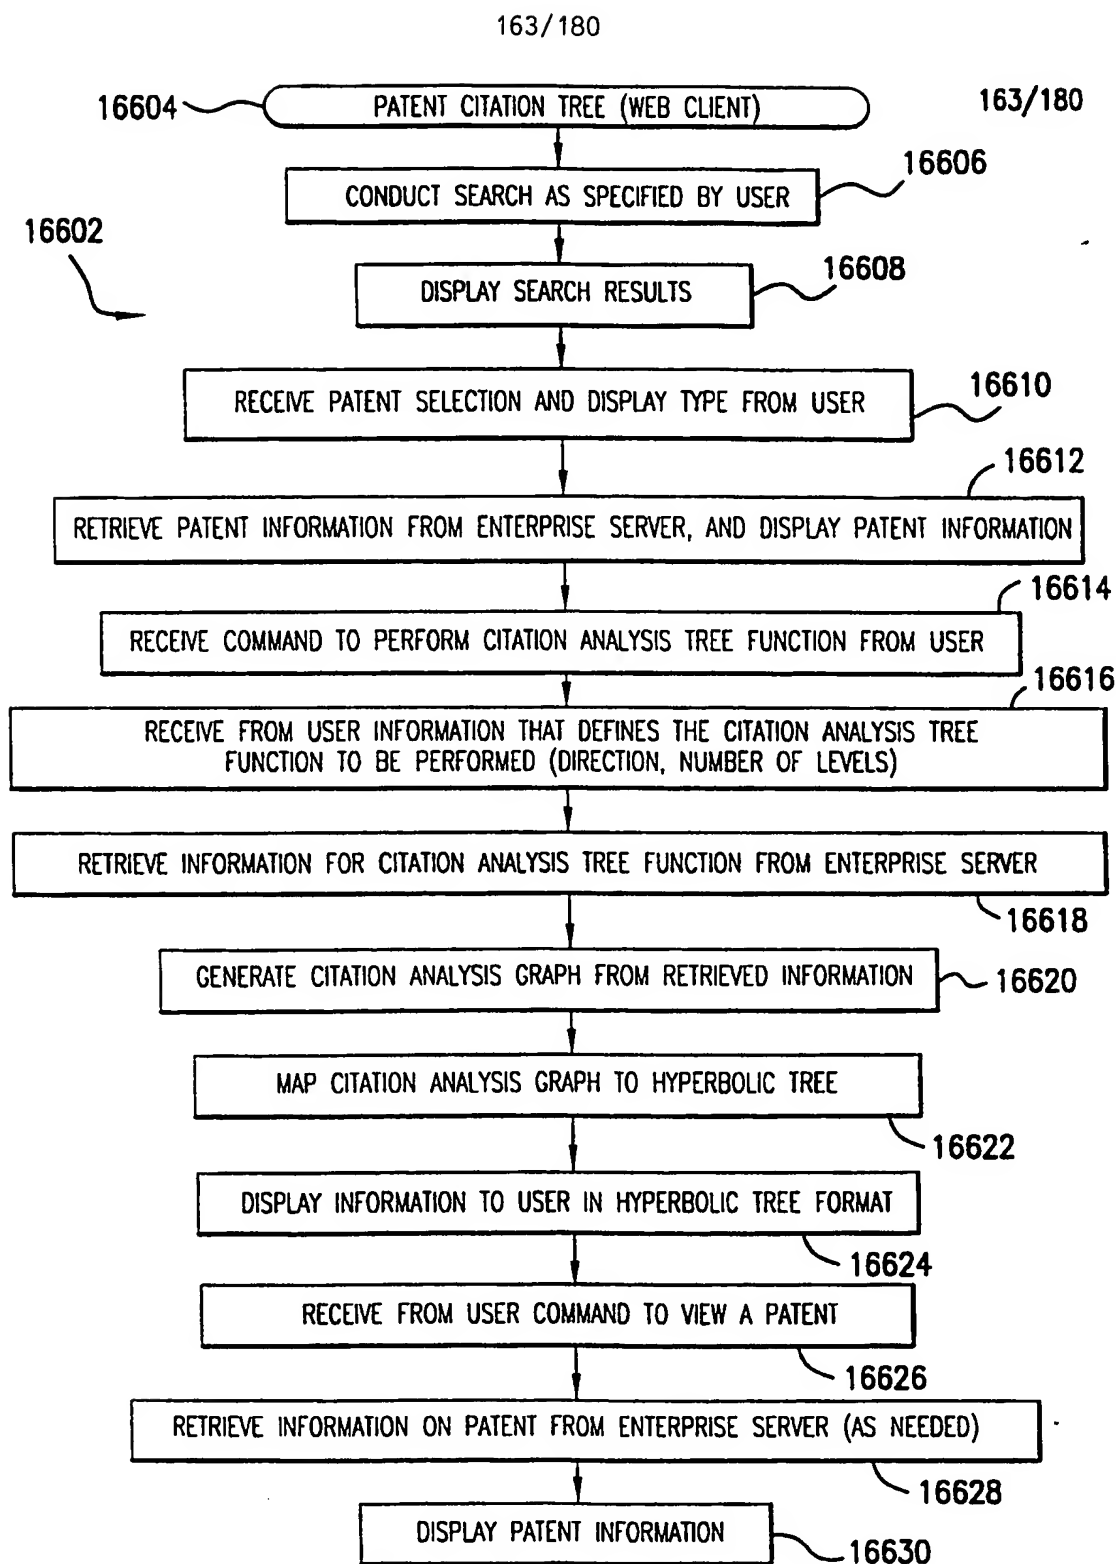


FIG. 166

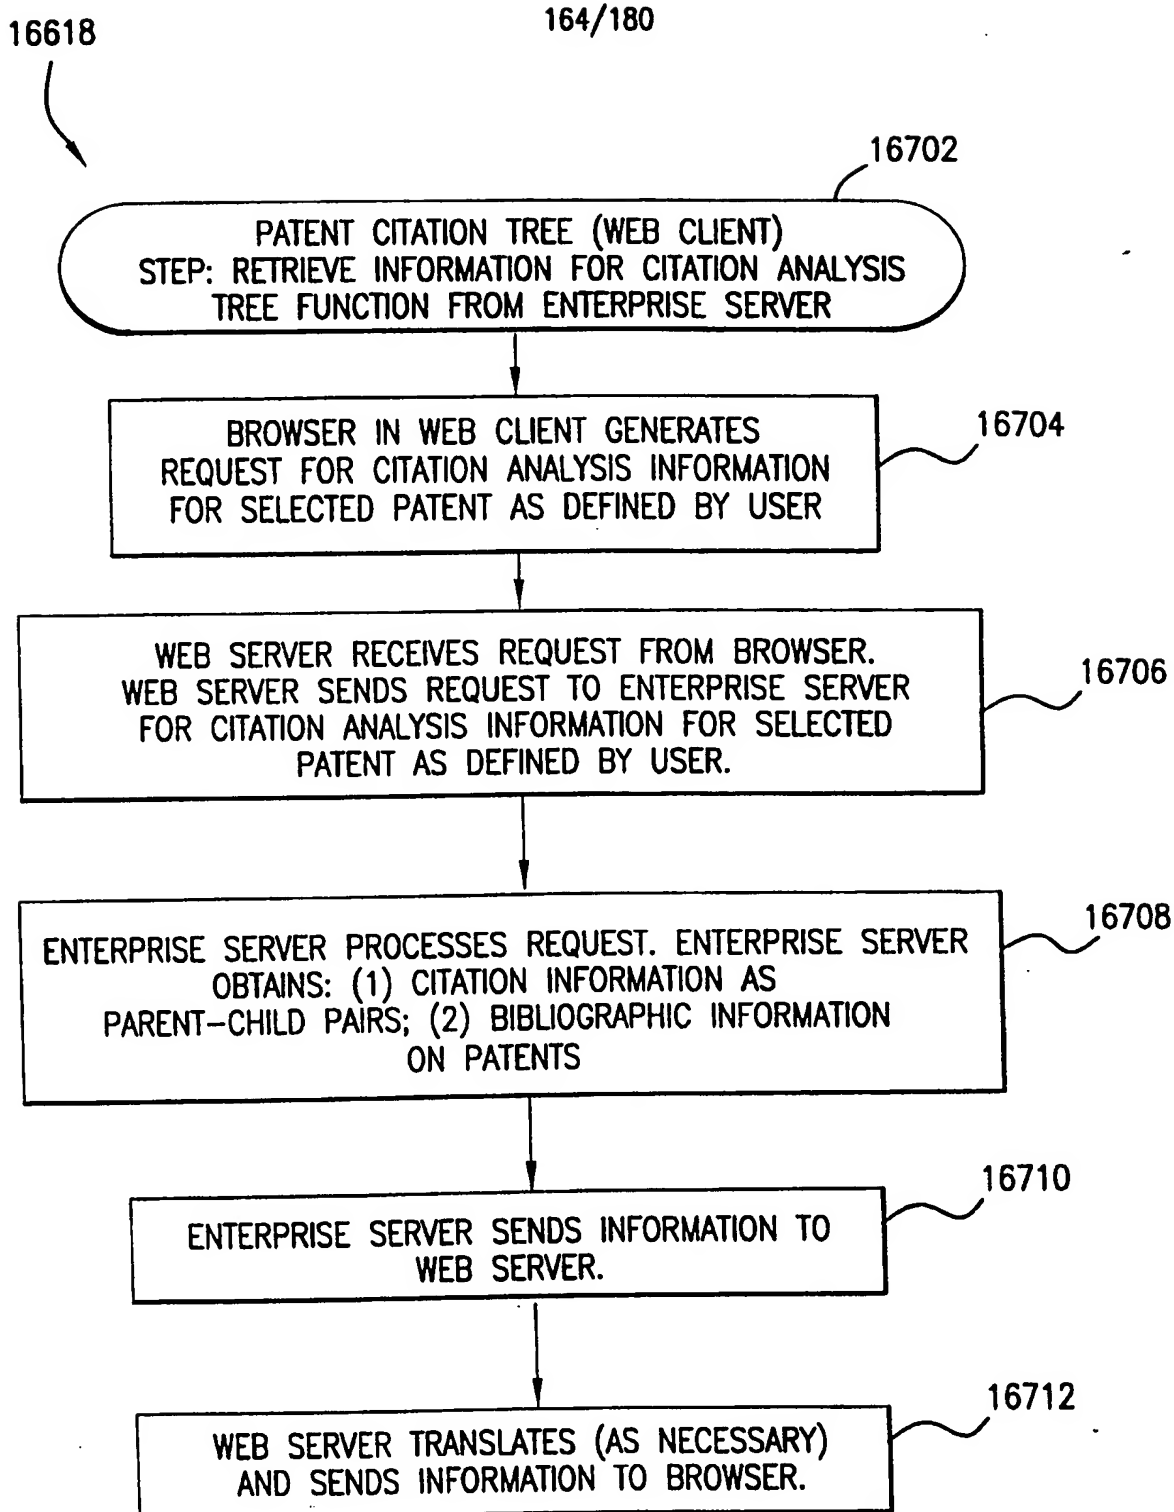
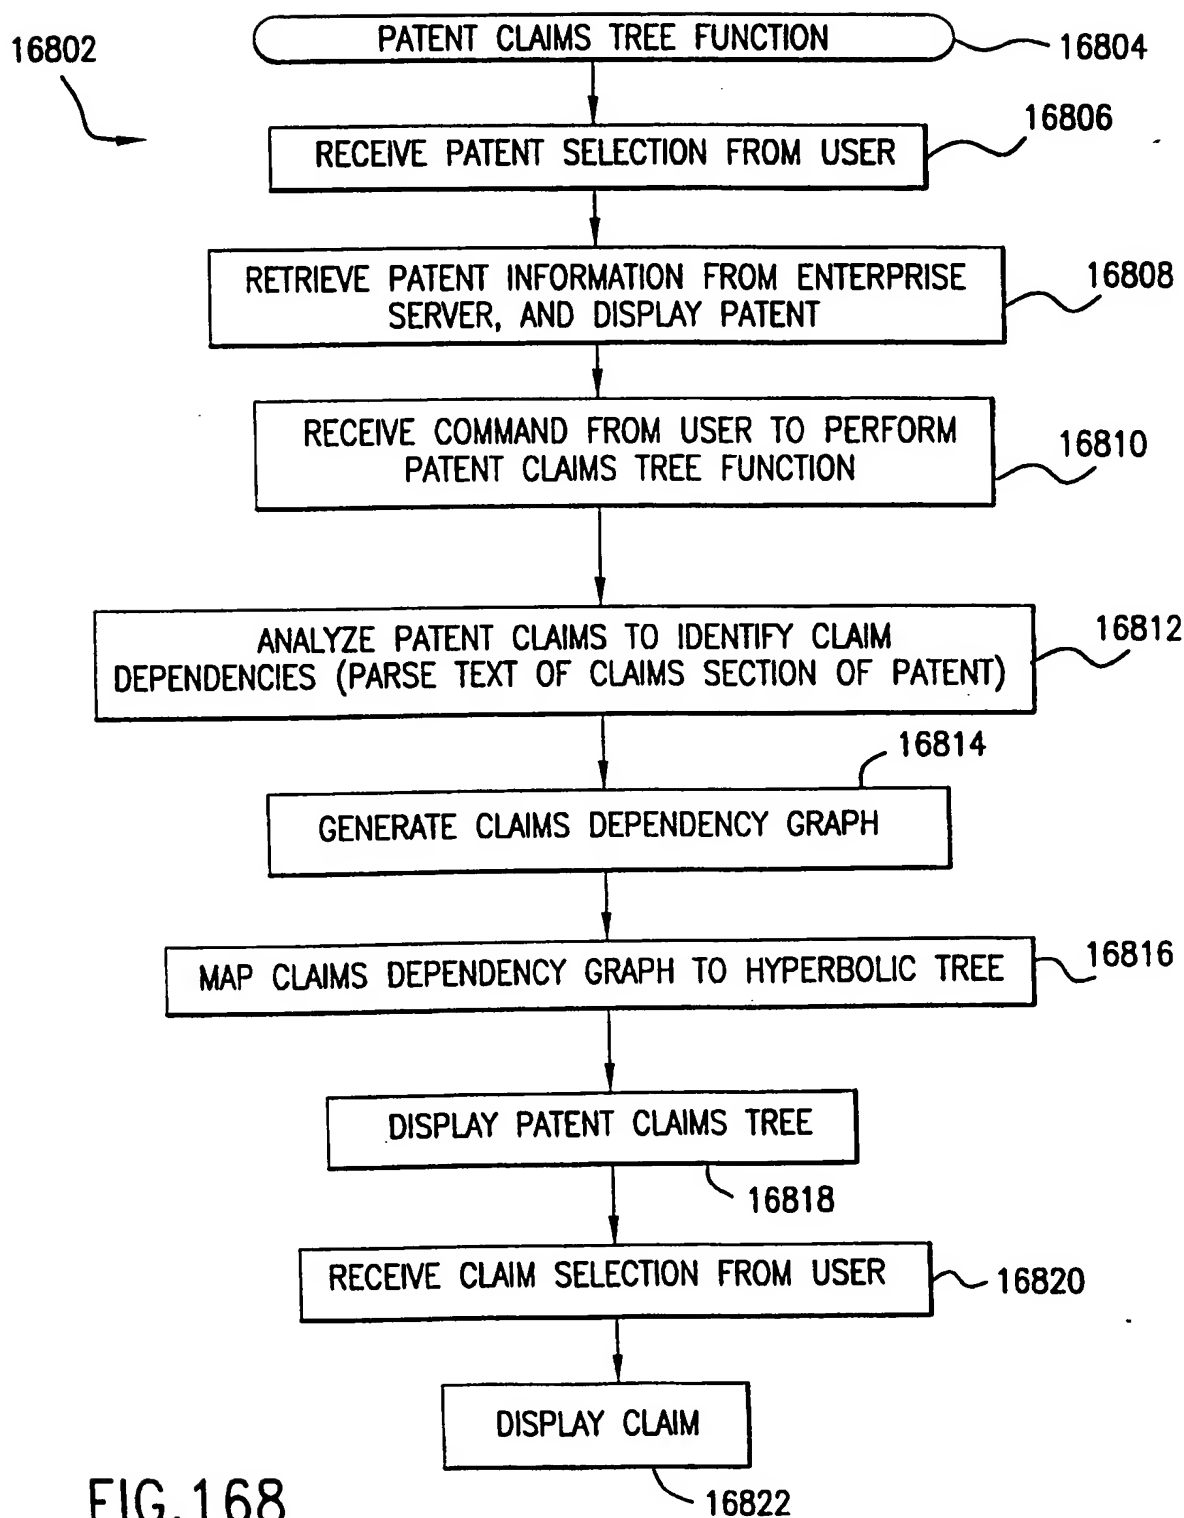


FIG.167

165/180



166/180

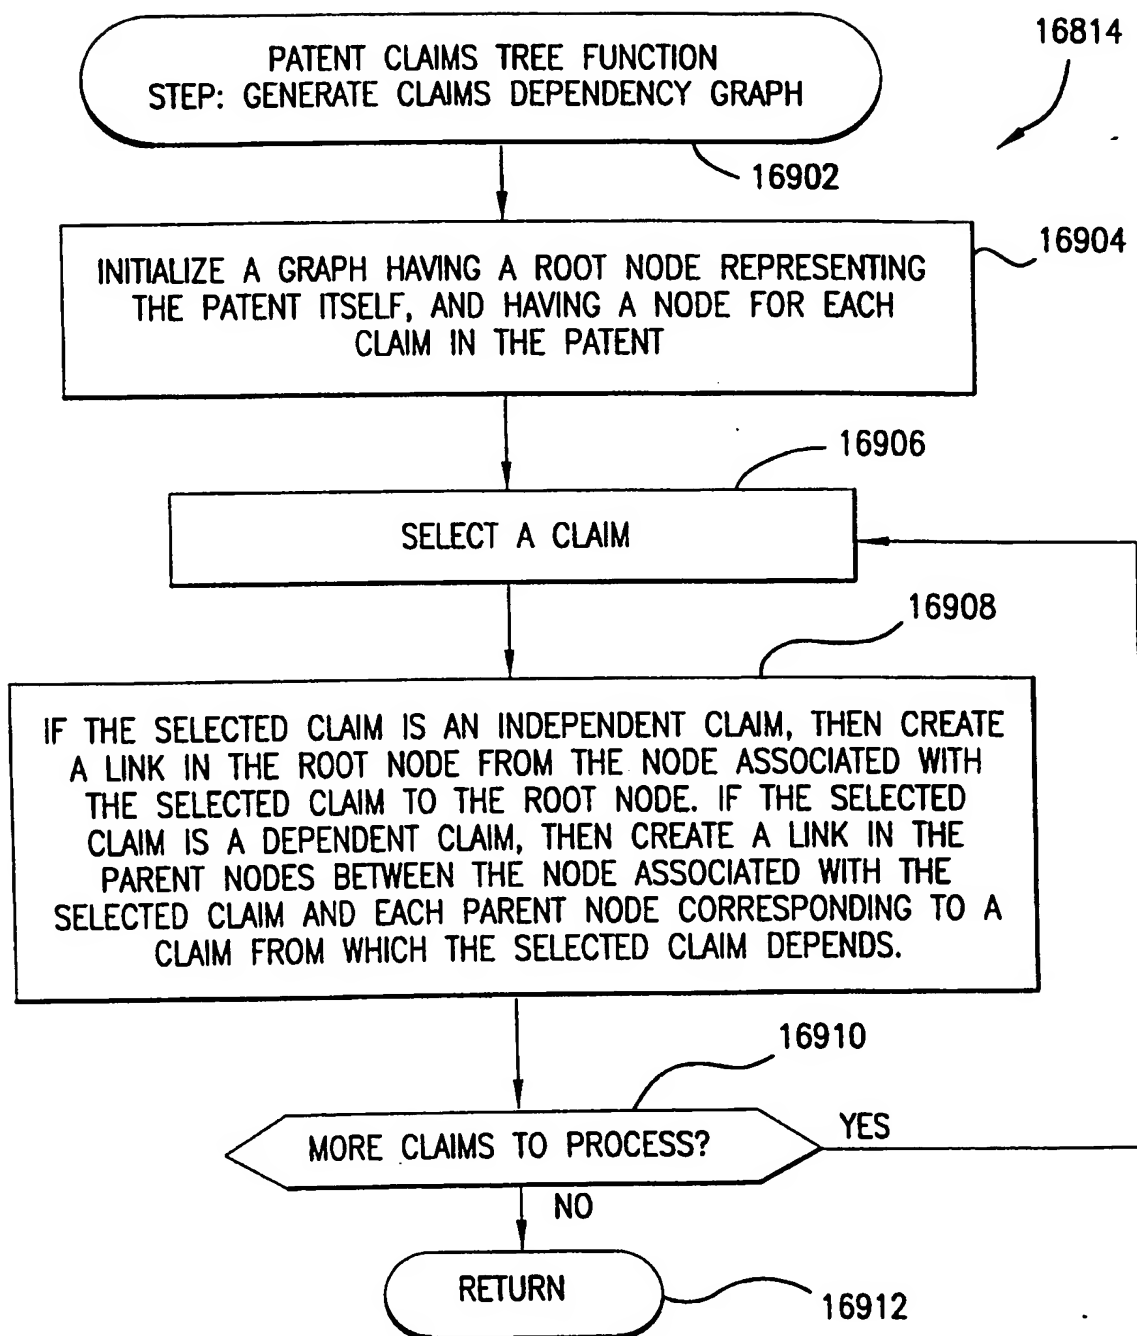
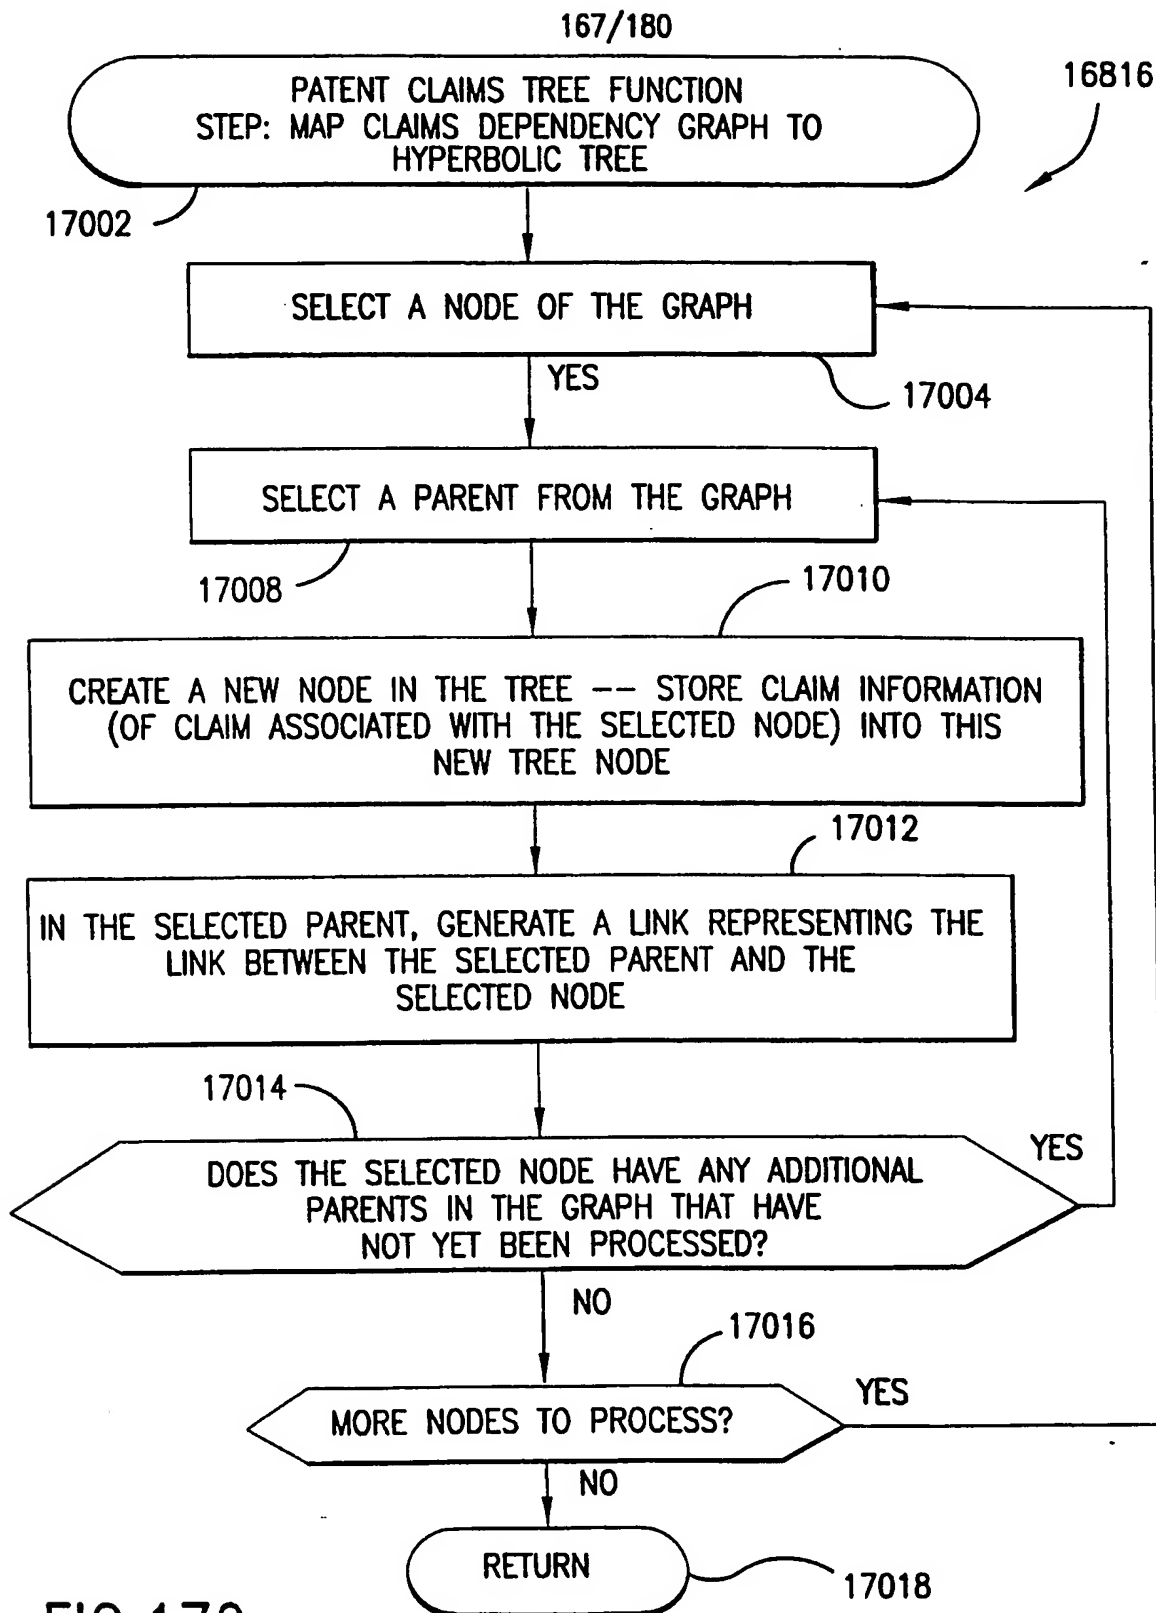


FIG.169



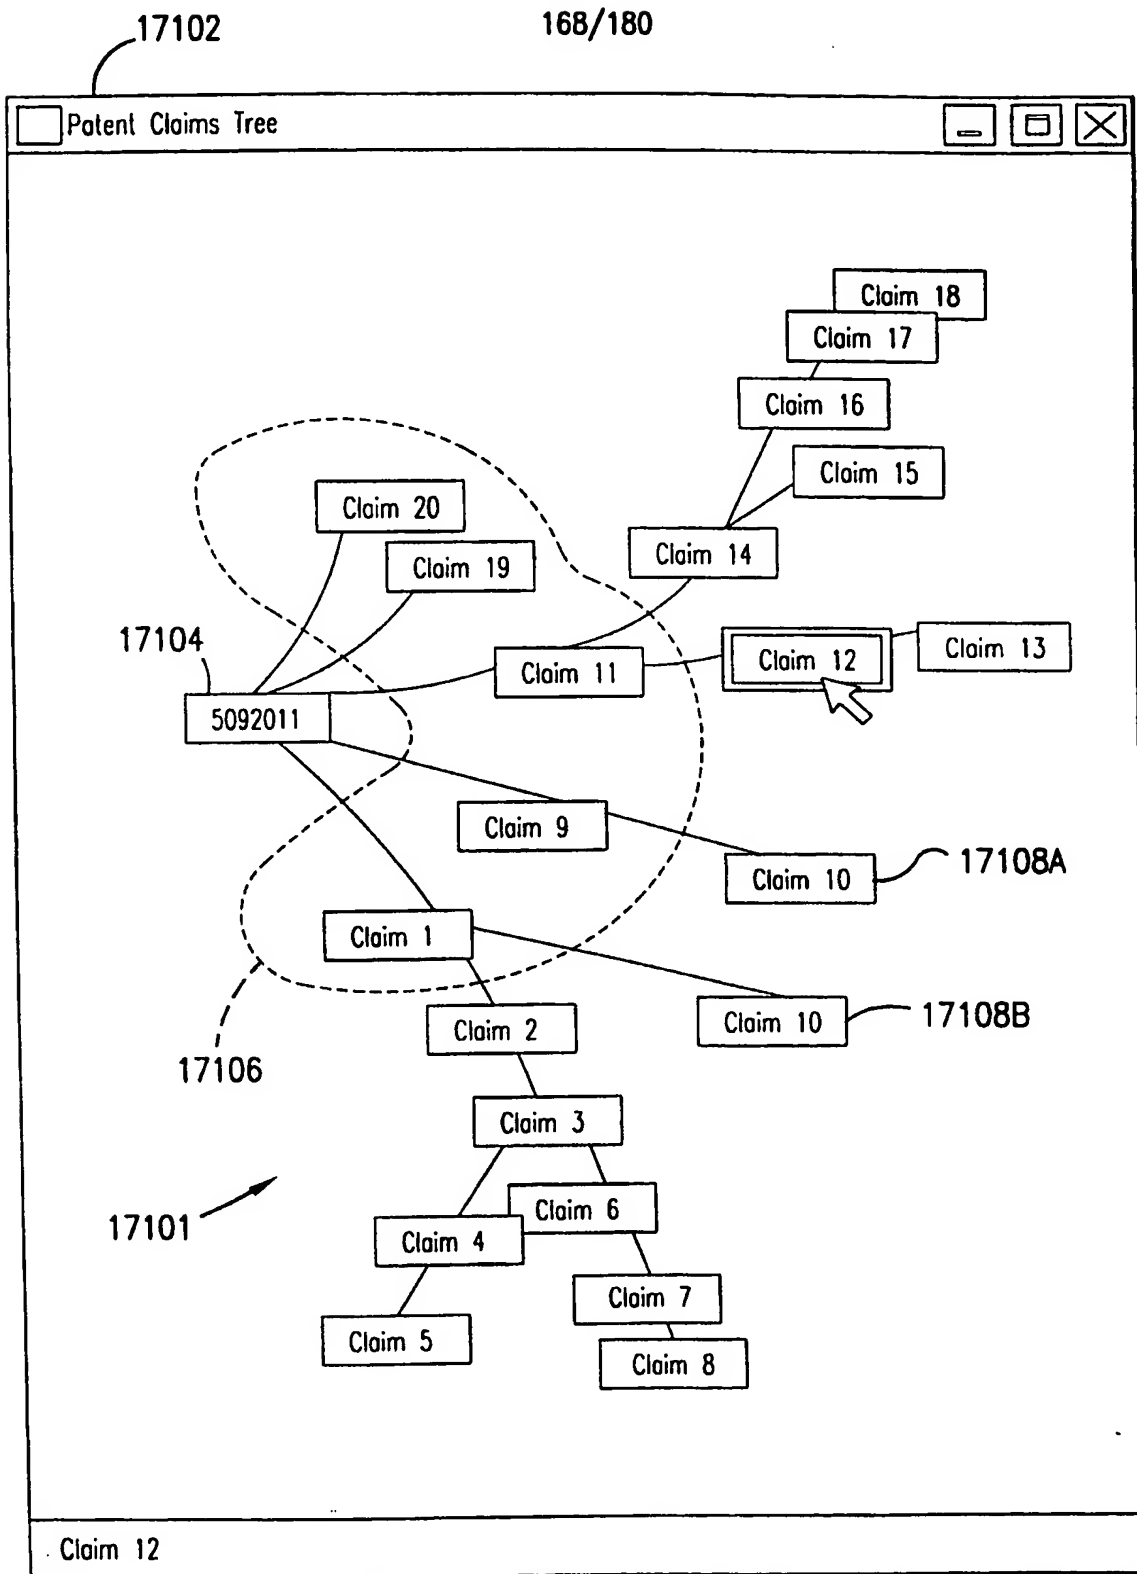


FIG. 171

169/180




FIG.172

170/180

17302




FIG.173

171/180




FIG. 174




FIG. 175

172/180

15718

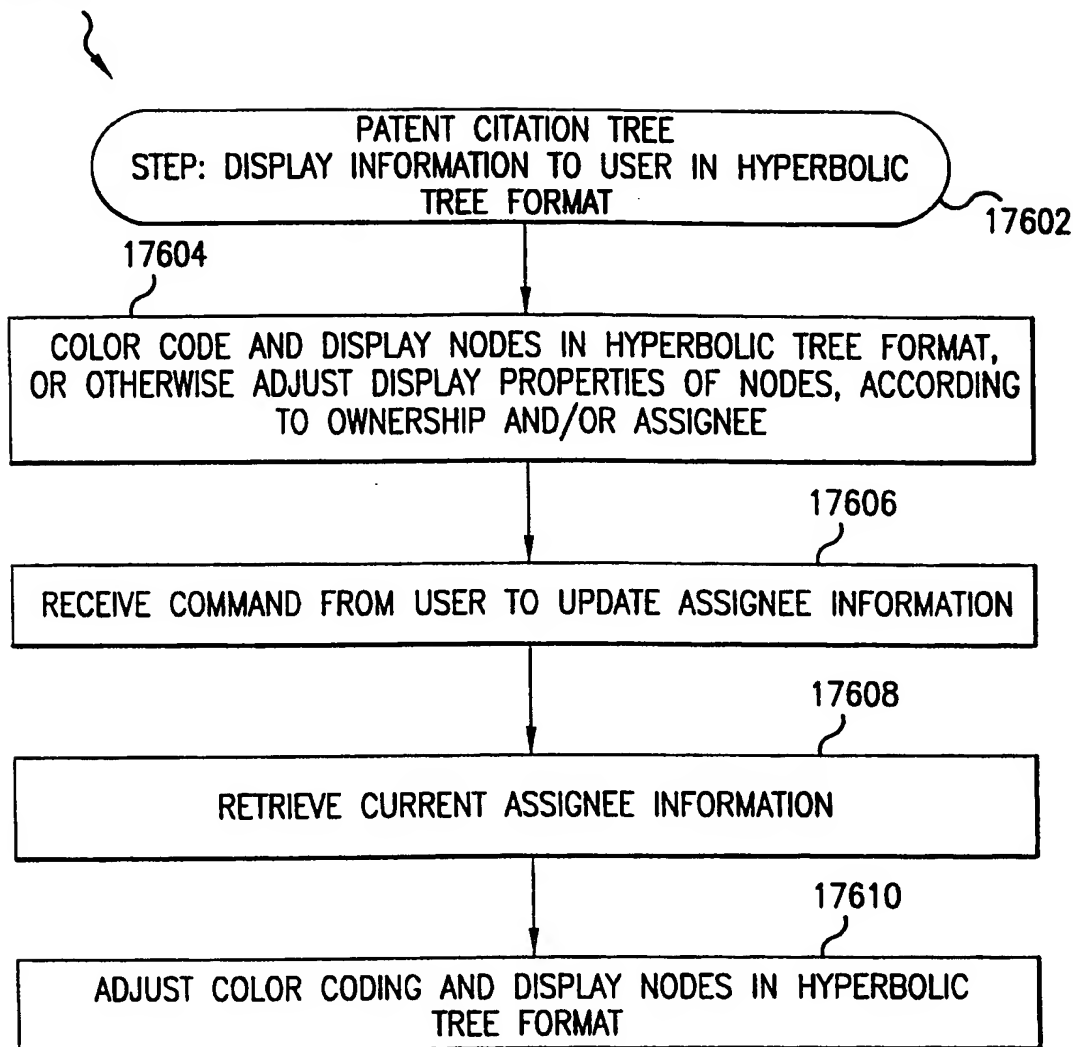


FIG.176

173/180




FIG.177

174/180

WINDOW
17702




FIG.178

175/180

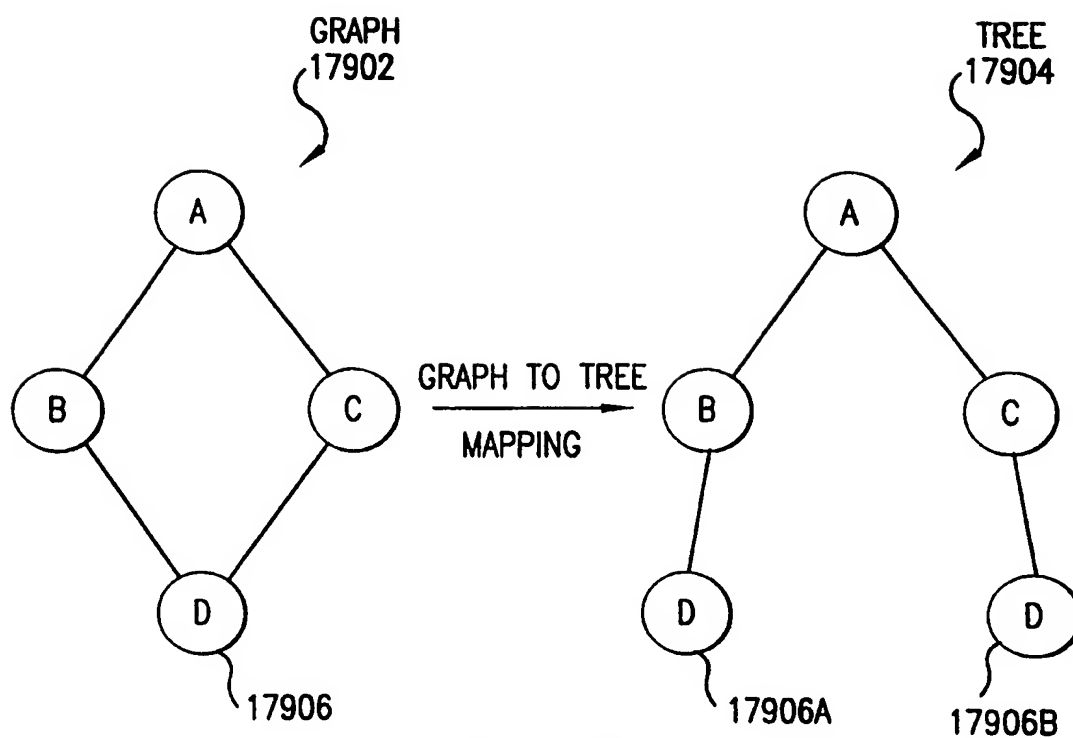


FIG.179

176/180

PARENT/CHILD TABLE
18002

PARENT	CHILD
1	2
1	3
2	4
4	5
4	6
5	7
6	7

FIG.180


CITATION
ANALYSIS
GRAPH
18102

FIG.181

177/180

PATENT BIBLIOGRAPHIC
INFORMATION TABLE

18202

PATENT	BIBLIOGRAPHIC INFORMATION
1	BIB INFO FOR PATENT 1
2	BIB INFO FOR PATENT 2
3	BIB INFO FOR PATENT 3
4	BIB INFO FOR PATENT 4
5	BIB INFO FOR PATENT 5
6	BIB INFO FOR PATENT 6
7	BIB INFO FOR PATENT 7

FIG.182


CITATION
ANALYSIS
TREE
18302

FIG.183

178/180

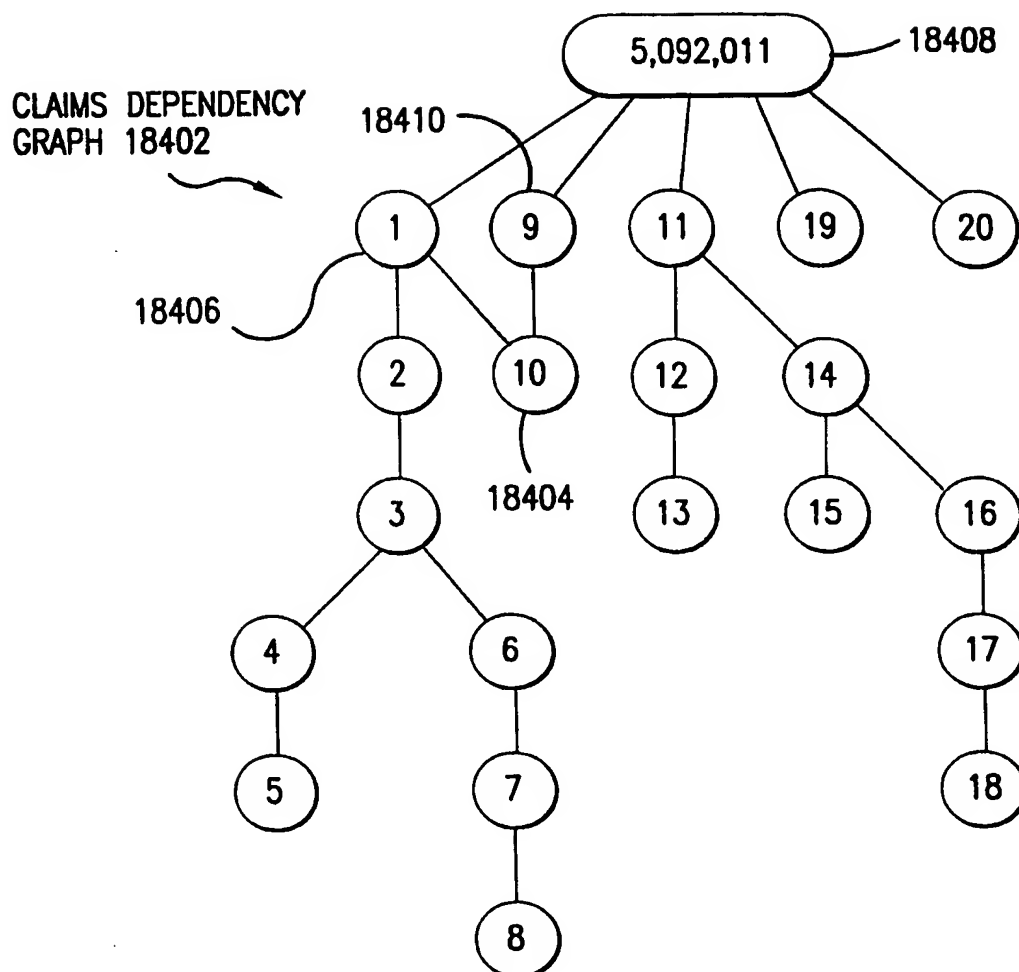


FIG.184

179/180


CLAIMS DEPENDENCY
TREE 18502

FIG.185

180/180




FIG.186

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/10923

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 623 679 A (RIVETTE KEVIN G ET AL) 22 April 1997 see abstract see column 4, line 50 - column 5, line 10 see column 10, line 8 - column 10, line 17 see column 14, line 27 - column 14, line 44 see column 27, line 9 - column 27, line 29 ---	1-82
A	US 5 544 352 A (EGGER DANIEL) 6 August 1996 see column 2, line 53 - column 8, line 2 see figures 4A, 4B -----	1, 27, 53, 79-82

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

12 October 1998

Date of mailing of the international search report

19/10/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Abbing, R

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US 98/10923

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5623679 A	22-04-1997	US 5623681 A	22-04-1997
		AU 5712796 A	07-11-1996
		CA 2218094 A	24-10-1996
		CN 1184539 A	10-06-1998
		EP 0832460 A	01-04-1998
		WO 9633465 A	24-10-1996
		US 5806079 A	08-09-1998
		US 5809318 A	15-09-1998
		US 5696963 A	09-12-1997
		US 5799325 A	25-08-1998
		AU 688836 B	19-03-1998
		AU 1292595 A	06-06-1995
		AU 7189998 A	13-08-1998
		BR 9408111 A	05-08-1997
		CA 2176729 A	26-05-1995
		CN 1141093 A	22-01-1997
		DE 731948 T	15-05-1997
		EP 0731948 A	18-09-1996
		JP 9505422 T	27-05-1997
		WO 9514280 A	26-05-1995
US 5544352 A	06-08-1996	AU 7207494 A	17-01-1995
		CA 2164954 A	05-01-1995
		EP 0704075 A	03-04-1996
		WO 9500896 A	05-01-1995

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☒ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:**

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.